New Results on the Predictive Value of Crude Oil for US Stock Returns

Matthew Brigida*

September 18, 2015

Abstract

In an original analysis, Driesprong, Jacobsen, and Maat (2008) found that crude oil is able to predict stock market returns. They further find evidence that this predictive relationship is a true market inefficiency, driven by slow information diffusion, and not simply due to time-varying risk premia. While confirming these previous results, we find that the nature of the predictive relationship changed in the latter half of 2008. After mid-2008, the predictive relationship switched signs, and exhibited characteristics which make it much more likely that the predictive relationship is due to time-varying risk premia rather than a market inefficiency.

1 Introduction

Crude oil is the largest commodity market, and the world consumption of 70–80 million barrels per day drives economic growth. Accordingly there has been a great deal of research on the link between crude oil price changes and the macroeconomy. There have been somewhat fewer analyses on the relationship between crude oil changes and stock market prices (prominent among these are Chen, Roll, and Ross (1986) and Jones and Kaul (1996)).

However until Driesprong, Jacobsen, and Maat (2008) (hereafter DJM) there had been no research on whether crude oil price changes affected future changes in the stock market. This is despite the two plausible scenarios that, firstly, a shock to crude oil prices increases (time-varying) risk premia which then affects

*Associate Professor of Finance, Clarion University of Pennsylvania, 840 Wood St., Clarion, PA 16214, Email: mbrigida@clarion.edu or matt@complete-markets.com

stock returns. Secondly, crude oil price shocks may be slowly incorporated into stock prices, due to uncertainty regarding oil how oil affects certain stocks.

In their analysis, DJM found evidence for a negative relationship between crude oil returns and the following month's stock market returns. Given crude oil's ability to predict negative stock market returns, a lack of correlation between crude oil returns and other factors known to change with time-varying risk premia, and a somewhat greater predictive effect in non-oil related sectors, DJM conclude the evidence favors the predictive relationship being due to slow incorporation of information into prices.

In this analysis we also investigate the ability of crude oil to predict stock returns. In particular we build on and extend the work of DJM. We do so by allowing market participants to adapt to changes in the predictive relationship between crude oil and the stock market.

In their analysis DJM assumed a constant relationship between lagged oil returns and stock market returns. However, if the predictive relationship is due to slow information transmission between markets, then the parameters defining the relationship are likely to change. That is to say, market participants constantly attempt to improve their understanding of the interrelationships between markets. This learning process will be reflected in a time-varying coefficient linking past crude oil changes to present stock market changes. In this analysis we'll model the learning process of market participants through a time-varying parameter regression, where the parameters are estimated using the Kalman filter.

Our results will also have implications for the mechanism by which crude oil returns are able to predict stock returns. DJM conclude the predictive ability of oil is not due to time-varying risk premia. Instead they find it is due to a market inefficiency. They hypothesize that oil predicts future stock returns through the gradual information diffusion hypothesis of Hong and Stein (1999). Investors consistently underreact to oil price changes, and therefore act with a delay.

Their conclusion is in large part based on the fact that, over their sample which ended in 2003, higher oil prices predict lower stock returns. The first contribution of the present analysis is to show that the relationship shifted to being strongly positive in 2008, and has remained positive since. We discuss the implications of the shift in the relationship on hypotheses for the underlying cause of the predictive relationship.

We also test for changing correlations between oil prices and macroeconomic variables which are known to predict time-varying risk premia. Specifically, we test for significant correlations between crude oil returns and the dividend yield, term spread, and default spread. We find, post-2008, crude oil prices became significantly correlated with both the dividend yield and the default spread.

If it is true that the oil's predictive ability is due to investors reacting to oil price changes with a delay, and not time-varying risk premia, then we should find oil's predictive ability to be contingent on the stock sector. That is, for example, a decrease in crude oil should cause an immediate decrease in oil services stocks. This is because the relationship between crude oil prices and the stocks are clear, and so investors should not act with a delay.

Conversely, if oil prices decline, and there is slow information diffusion, then we may see a delayed rise in non-oil related stocks. The rise in these stocks is delayed, because the link between crude oil prices and the stocks are unclear, and it takes time for market participants to grasp the implications of the crude decline for the market sector.

As a brief review, the following explains why crude oil could seem to predict future stock returns through time-varying risk premia. Consider stock prices following a process such as geometric Brownian motion $dS_t = \mu S_t dt + \sigma S_t dB_t$.

Then the absence of arbitrage implies there exists a vector λ such that $\mu - r = \sigma\lambda'$ where λ is the vector price of risk, r is the risk free rate, and μ and σ are vectors of expected returns and volatilties respectively. Thus changes in asset expected returns may be explained by changes in r, σ or λ. Prior research (Fama and French (1989)) has posited that the dividend yield, term spread, and default spread predict stock returns because these variables contain unique information about future business conditions.

The remainder of the paper is organized as follows. Section 2 describes the data used in the analysis, section 3 the methods and results, and section 4 concludes.

2 Data

Throughout this analysis we use returns on the S&P 500 index and Select Sector SPDR ETFs. We use returns on these securities because they can be traded easily and with little cost. This is opposed to other academic indices, which while often very useful for understanding relationships between economic and financial variables, cannot be easily traded. Since ultimately this analysis is concerned with a predictive relationship, which is most useful if it can be used in low-cost trading strategies, we prefer to use tradeable instruments.

Oil prices are continuous front-month NYMEX West Texas Intermediate (WTI) crude oil futures (ticker: CL). This contract is for delivery in Cushing, OK. For robustness, we also use continuous front-month ICE Brent crude oil futures (ICE ticker: B). Brent crude oil is deliverable to the Sullom Voe terminal in the European North Sea. All stock and crude oil data are aggregated into monthly continuously compounded return series.

Note, in crude oil markets the one-month (or front-month) futures contract is often used as a quote for the spot price. This is because, pipeline nominations are made at the end of each month, for the following month. This is where you schedule your crude oil's passage through the pipeline network. So given the time it takes to arrange shipment, spot transactions are really for for delivery next month. The one-month futures contract therefore is nearly identical to the spot price, with the benefit of slightly clearer reporting.

3 Method & Results

We first estimate the predictive regression using monthly continuous front-month WTI crude oil futures returns, and returns on the S&P 500 value weighted stock market index.

$$r_t^m = \alpha_0 + \alpha_1 r_{t-1}^{oil} + \mu_t \quad \text{where} \quad \mu_t = r_t^m - E_{t-1}(r_t^m) \tag{1}$$

where r_t^m and r_t^{oil} denote continuously compounded returns on the S&P 500 and crude oil respectively, α_0 is a constant term, and μ_t is the usual error term. In estimating the equation over our full-sample we find no evidence for the ability of crude oil returns to predict stock returns (coefficient is -0.0044 with a -0.1590 t-value).

This result contradicts that in DJM (their t-value was -3.47 on a similar regression of monthly US value-weighted markets index returns on lagged WTI returns). However their sample ranged from October 1973 to April 2004, and our sample is from May 1983 to November 2014, so this is possible the source of the conflicting results.

To investigate we restrict our sample to May 1983 through April 2004, and our find parameter estimates (-0.0856 coefficient with a -2.461 t-value) agree with DJM. Estimating the same equation on the period May 2004 to November

2014 afford a coefficient of 0.1439 with a 3.3510 t-value. These results imply there was a marked shift in the predictive relationship between the two sample periods.

To investigate the timing and nature of the shift in the relationship we translate the predictive regression into state-space form, to allow for time-varying coefficients. The model parameters are estimated via the Kalman filter, and prediction-error decomposition.

Measurement Equation

$$r_t^m = \alpha_{0,t} + \alpha_{1,t} r_{t-1}^{oil} + \mu_t, \quad \mu_t \sim N(0, \sigma_\mu^2)$$

Transition Equations

$$\alpha_{0,t} = \xi_0 + \eta_0 \alpha_{0,t-1} + \upsilon_{0,t}, \quad \upsilon_{0,t} \sim N(0, \sigma_{\upsilon_0}^2)$$

$$\alpha_{1,t} = \xi_1 + \eta_1 \alpha_{1,t-1} + \upsilon_{1,t}, \quad \upsilon_{1,t} \sim N(0, \sigma_{\upsilon_1}^2)$$

where $Cov(\mu_t, \upsilon_{0,t}) = Cov(\mu_t, \upsilon_{1,t}) = Cov(\upsilon_{1,t}, \upsilon_{0,t}) = 0$. In this formulation of the predictive relationship, the parameters linking crude oil in month $t-1$ to stock market returns in month t are allowed to evolve according to an autoregressive process. This allows the model to incorporate learning by market participants. That is, within the model market participants are allowed to become aware of the predictive relationship over time, which would lead to a lessening of the predictive relationship ($\alpha_{1,t} \to 0$). It also allows a time-varying risk premia predictive effect to dominate the slow information diffusion effect (in which case we would expect $\alpha_{1,t} > 0$ for $t > t^*$).

Figure 1 below shows the filtered coefficient estimate. We immediately see a marked increase in the coefficient corresponding to September 2008. It therefore appears that the 2008 financial crises had a significant effect on the predictive relationship between crude oil and stock market returns. This result, in itself, is of practical importance in applying the results of DJM, but it also highlights that we must consider the sample periods separately in the analysis which follows. Note, the results above are also found when using monthly Brent North Sea one-month crude oil futures prices.

Figure 1: Time-varying slope coefficient in the predictive regression.

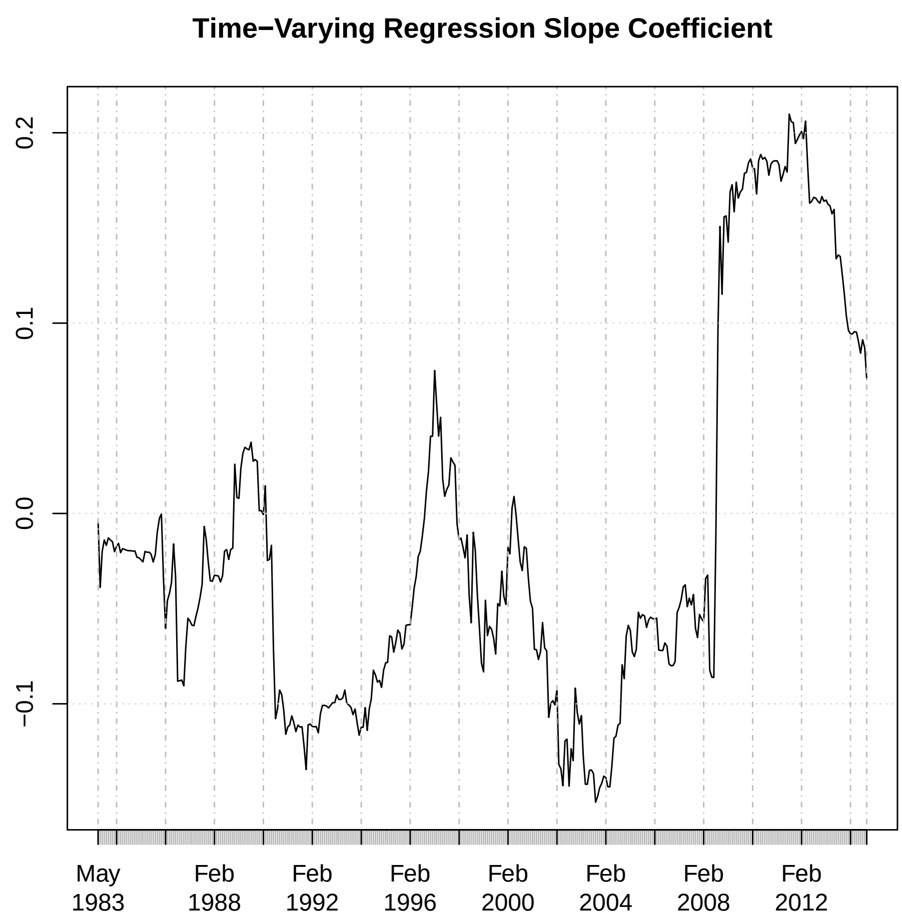

3.1 Agreement with Equilibrium Theory

At this point we can note that the positive coefficient post-2008 is in agreement with the theoretical predictive relationship between crude oil and the stock market. Hamilton (2003) shows oil price increases should lead to increased risk in the economy. This, in turn, should lead to higher future average stock market returns.

So this evidence is consistent with oil predicting movements in the stock market due to time-varying risk premia, and not a market inefficiency. Crude oil movements affect market risk, which is then priced into equity returns.

3.1.1 Why the shift in the coefficient in 2008?

Crude oil's coefficient in the predictive regression increased sharply from -0.09 in July 2008, to 0.16 in October 2008. At first glance, this would seem to certainly be caused by the 2008 financial crisis.

However, increase in the coefficient coincides perfectly with the lowest point of the US production of crude oil since 1946[1] The increase in production after this low point is largely driven by light oil from tight formations such as the Bakken Shale in western North Dakota. This low point thus represents the point at which the market certainly recognized the shale revolution in North America. Therefore, the cause of the shift in the predictive relationship may be somewhat more complicated than being simply an artifact of the financial crisis. This may be a fruitful avenue for further research, however in this analysis we will simply admit we cannot say for certain what caused the shift the predictive relationship.

[1] see: `http://www.eia.gov/beta/api/qb.cfm?category=296686&sdid=PET.MCRFPUS2.A`

Figure 2: Time-varying coefficient and US crude oil production. Coefficient values are on the left vertical axis, and US crude oil production (in millions of barrels per day) is on the right vertical axis.

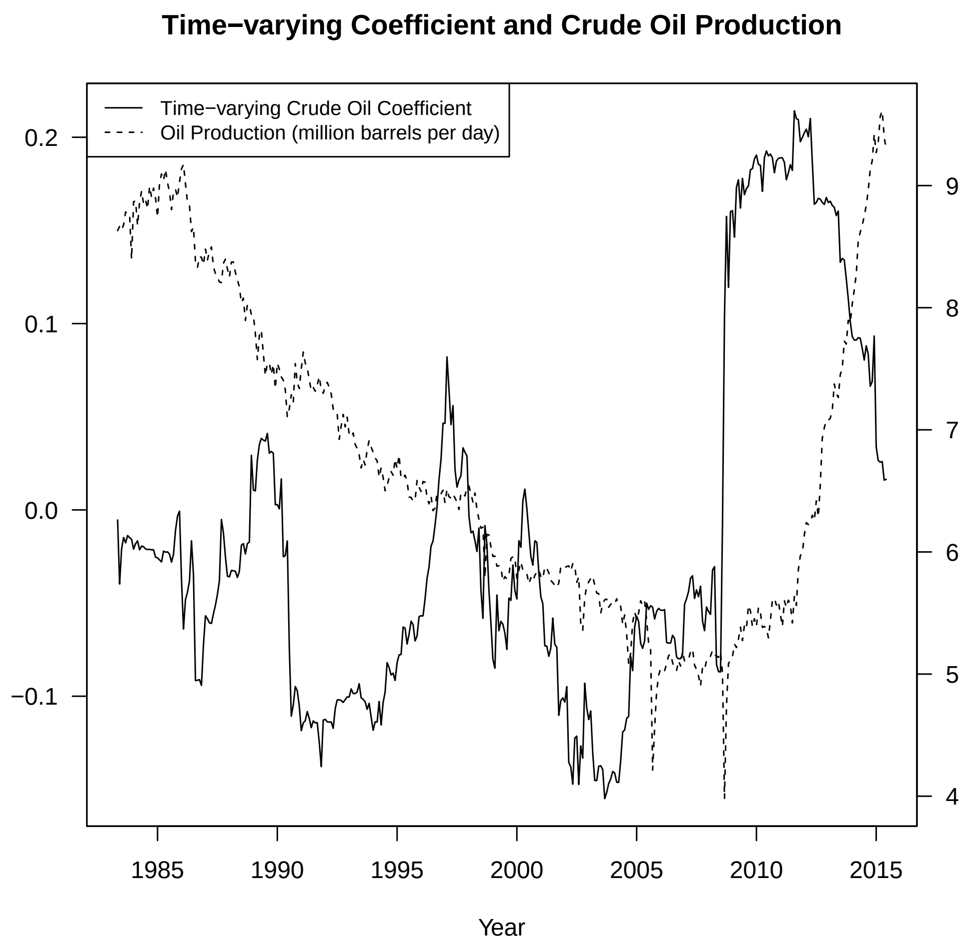

3.2 Predicting Negative Excess Returns

Schwert (2003) set a requirement that a true market efficiency, such as slow information diffusion, must be able to predict negative excess returns. This is a fairly stringent requirement, but pre-2004 DJM found that this requirement was met. However, the results of our time-varying parameter model above have shown an abrupt shift in the predictive regression parameters in August-September 2008. Our question is therefore, has the ability for crude oil returns to predict negative market excess returns persisted?

To test this we estimate a linear predictive regression over the sample September 2008 to May 2015. Doing so affords the following parameters:

$$E_{t-1}\left(r_t^m\right) = 0.0084 + 0.1566 r_{t-1}^{oil} \tag{2}$$

The monthly standard deviation of oil price changes is 9.64%, an oil price decline 1 standard deviation below the -0.8146% mean return would predict a negative excess return on the market so long as the monthly risk-free rate is below 0.7971% (9.5663% annually). Since the risk-free rate was substantially below this over the sample period, we conclude oil price changes can predict negative excess market returns.

In fact, for any monthly oil return less than -5.0815%, the results predict a negative market excess return assuming the annual risk-free interest rate is 1%. Out of the 82 months in the sample period, 21 months had an oil return less than -5.0815%, and 10 months had an oil return less than -10.4546% (one standard deviation the average oil return). So we conclude oil price changes do predict negative excess market returns.

3.3 Including Longer Lags of Oil

Further evidence was found in favor of information diffusion by DJM when they included lagged changes in oil in their predictive regression. The coefficients on lagged oil were insignificant, implying oil's predictive ability is short-lived. This would not be the case if changes in oil prices acted on future stock market changes through their effect on time-varying risk premia.

We estimate the parameters of a purely predictive regression $r_t^m = \alpha_0 + \sum_{i=1}^{6} r_{t-i}^{oil} + \mu$, where m and oil denote monthly market and oil returns respectively, as in DJM. We do so over the post-crisis 2009 to 2013 period. In this model the lags of oil returns are significant out 5 months. This result points to the predictive relationship being driven by time-varying risk premia.

3.4 Oil and Macroeconomic Factors

In this section we'll test whether the relationship between oil and factors, which are known to predict stock time-varying risk premia, changed after the 2004 data cutoff date used in DJM. We'll further divide the data set into windows around the 2008 financial crisis. This is because of the marked change in the predictive relationship in 2008 as shown by our time-varying parameter model, as well as the obvious potential effect a financial crisis may have on estimated correlation coefficients. This section is informative because the lack of correlation prior to 2004, found in DJM, was cited as evidence that the predictive ability of oil was not due to an effect oil had on macroeconomic factors.

Testing for significant correlation between crude oil price changes and the term spread on weekly data ranging from May 1983 to April 2004 finds an insignificant 0.0134 coefficient. From May 2004 to December 2007 the coefficient is -0.1827 and insignificant. The coefficient estimated over the period January 2008 to May 2015 is 0.0544 and is also insignificant.

Table 1: Summary of correlations of oil price changes with macroeconomic variables. P-values are below the correlation coefficients in parentheses. ****, ***, **, * denote significance at the 0.1%, 1%, 5%, and 10% levels respectively.

Macro Variable	Correlation Coefficient		
	5/1983–4/2004	5/2004–12/2007	1/2008–5/2015
Default Spread	-0.0480	-0.5433	-0.2836
	(0.4404)	(0.0000)****	(0.0130)**
Term Spread	0.0134	-0.1827	0.0544
	(0.8293)	(0.1622)	(0.6404)
Dividend Yield	-0.0842	-0.6085	-0.2385
	0.1757	(0.0000)****	(0.0379)**

The default spread's correlation with oil price changes in -0.0480 and insignificant from April 1983 to April 2004. However, from May 2004 to December 2007 the coefficient is -0.5433 and significant at the 0.1% level. Over the January 2008 to May 2015 period the correlation is -0.2836 and significant at the 5% level.

The dividend yield's correlation is -0.0824 and insignificant from April 1983 to April 2004. From May 2004 to December 2007 the correlation coefficient is -0.6085 and significant at the 0.1% level. Over the January 2008 to May 2015 period the correlation is -0.2385 and significant at the 5% level.

An insignificant correlation is consistent with the predictive relationship being caused by slow information diffusion and not time-varying risk premia. Over the time period used in DJM, our result of insignificant correlation coefficients for all three economic variables is consistent with DJM. However, two of the three variables are significant from May 2004 to May 2015. Moreover, this is not simply an artifact of the 2008 financial crises. When subdividing this period into subintervals from May 2004 to December 2007, and January 2008 to May 2015, the two variables are negative and significant over both subintervals.

On balance, the evidence is consistent with DJM that pre-2004 the predictive relationship may have been due to slow information diffusion. However post-

2004 our evidence points to time-varying risk premia as being a possible cause of the predictive relationship.

3.5 Sector Analysis

Given the results above, we will investigate crude oil's ability to predict the following month's stock sector return for the periods before and after September 2008. DJM found crude oil's predictive ability is greater in non-oil related industries. This finding is consistent with their hypothesis of slow information diffusion. That is, market participants can easily gauge the impact of oil price changes in sectors of the economy intimately involved with crude oil production or consumption. So in these oil related sectors, oil should have little predictive ability—information is incorporated immediately. However, stock in sectors unrelated to crude oil should react more slowly to oil price changes, because the oil's effect is more uncertain, thereby creating the predictive relationship. Alternatively, if oil's predictive ability spans all stock sectors, particularly including the energy production sector, then this is more likely do to time-varying risk premia.

The results of our sector analysis are in table 2. Over the full sample we get results similar to DJM. Oil changes are able to predict returns in the financial, materials, and industrial sectors only. The uncertain transmission mechanism from oil price changes to financial stock returns would support slow information diffusion. Similarly, crude oil might have an unclear effect on materials sector (mining, steel making) and the general industrial sector. Importantly, oil does not predict returns in the energy sector. In our pre-2008 sample we find no predictive relationships.

In our post-2008 sample, there is a predictive relationship for every stock sector other than health insurers and utilities. The predictive relationship is even present for the energy production sector. This result is evidence that crude oil affects the stock market through time-varying risk premia rather than slow information diffusion.

Table 2: Summary of ETF market model betas. The full sample period is from February 1999 to December 2014. Subsample 1 ranges from February 1999 to August 2008. Subsample 2 ranges from September 2008 to December 2014. Returns are monthly and continuously-compounded. All coefficients are significant at the 0.1% level.

ETF Ticker	Short Description	Market Model β		
		Full-sample	Sub-sample 1	Sub-sample 2
XLV	Health Insurers	0.7345	0.7524	0.7007
XLE	Energy	0.9119	0.7193	1.0997
XLP	Consumer Staples	0.4898	0.3606	0.6379
XLK	Technology	1.3834	1.799	0.9415
XLF	Financial	1.2256	0.9791	1.5351
XLB	Materials	1.1720	1.0288	1.3126
XLU	Utilities	0.4918	0.4752	0.4652
XLY	Consumer Discretionary	1.0824	1.0358	1.1327
XLI	Industrials	1.1296	1.0116	1.2491

Table 3: Summary of the relationship between lagged crude oil returns and ETF returns. The full sample period is from February 1999 to December 2014. Subsample 1 ranges from February 1999 to August 2008. Subsample 2 ranges from September 2008 to December 2014. Returns are monthly and continuously-compounded. *, **, *** and **** denote significance at the 10%, 5%, 1% and 0.1% levels respectively.

ETF Ticker	Short Description	Coefficient on lagged crude oil returns		
		Full-sample	Sub-sample 1	Sub-sample 2
SPY	S&P 500	0.0638	-0.0751	0.1694***
XLV	Health Insurers	0.0311	-0.0189	0.0568
XLE	Energy	0.0905	-0.0473	0.1528 **
XLP	Consumer Staples	0.0485	-0.0175	0.0976**
XLK	Technology	0.0560	-0.0996	0.1689***
XLF	Financial	0.1703***	-0.0833	0.3696****
XLB	Materials	0.1205**	-0.0773	0.2342***
XLU	Utilities	0.0286	-0.0244	0.0333
XLY	Consumer Discretionary	0.0729	-0.1277*	0.2306****
XLI	Industrials	0.0950**	-0.1019*	0.2310****

4 Discussion & Conclusion

Using data up to April 2004, our results in this analysis confirm those found earlier by DJM. However, on balance our evidence finds, that after April 2004, the predictive relationship between crude oil and the stock market has likely been driven by time-varying risk premia.

Our evidence centers on four main findings. Firstly, at some time after April 2004 (or evidence points to about August-September 2008) there was a dramatic shift in the predictive relationship between crude oil and stock market. The coefficient of lagged oil returns went from about -0.1 to over 0.1 in a matter of months. This positive relationship agrees with the equilibrium theory that oil price increases, increase market risk premia, thereby increasing future stock market returns.

Secondly, oil lags out 5 months predict stock market returns in the post-2008 period. Thirdly, after April 2004, two of three economic variables well-known to forecast stock returns (the default spread and dividend yield) became correlated with crude oil.

Lastly, after April 2004 crude oil does not affect stocks in the energy sector any more quickly than stock in other sectors. This would imply information is not diffusing slowly through stock sectors, but rather is affecting stock prices through a common mechanism over across all stocks. This results makes particular sense given algorithmic trading in the latter half of the 1990s, and the general increases in speed with which market participants now react to information.

As pointed out by Schwert (2003), an extreme requirement for a market inefficiency such as slow information diffusion, is that market excess returns should be predictably negative. As we have shown, post-2004 the predictive relationship continued to predict negative excess market returns. This is evidence in favor of a market inefficiency.

In sum, our analysis finds compelling evidence that the predictive relationship between crude oil and stock market returns was dominated by the effect of time-varying risk premia post-2008. That is not to say it must continue to be so after our sample ends, but rather it is evidence of the competing forces of slow information diffusion and time-varying risk premia, and that one of these forces is not perpetually dominant of the other. Moreover, some of our evidence supported the market inefficiency claim (the ability to predict negative excess market returns).

References

Chen, N.F., R. Roll, and S. Ross, (1986). “Economic forces and the stock market.” *The Journal of Business* 59(3), 383–403.

Driesprong, G., B. Jacobsen, and B. Maat (2008). “Striking oil: Another puzzle?” *Journal of Financial Economics* 89, 307–327.

Fama, E.F., and K.R. French (1989). “Business conditions and expected returns on stocks and bonds.” *Journal of Financial Economics* 25, 23–49.

Hamilton, J.D., (2003). “What is an oil shock?” *Journal of Econometrics* 113(2), 363–398.

Hong, H.G. and J. Stein (1999) “A unified theory of underreaction, momentum trading and overreaction in asset markets.” *Journal of Finance* 54, 2143–2148.

Jones, C.M., and G. Kaul (1996). “Oil and the stock markets.” *The Journal of Finance* 51(2), 463–491.

Schwert, G.W. (2003) “Anomalies and market efficiency.” In: Constantinides, G.M., Harris, M., Stulz, R., (Eds.), The Handbook of the Economics of Finance, pp. 939–972.