

Commodity Prices and Industry Profitability*

Espen Henriksen[†]

Jens Sørlie Kværner[‡]

January 15, 2017

Abstract

Commodity-price changes may affect both the aggregate level of economic activity and the profitability of different industries. In this paper, we ask what is the relationship between commodity-price changes and long-run industry-level profitability, and answer this question through the lenses of equity-markets data and asset-pricing theory. We decompose the relationship between commodity-price shocks and equity-market valuations into changes in expectations of future cash flows and in revisions of discount rates, respectively. This methodology is applied to the oil market. Not surprisingly, we find that a negative oil-price shock is associated with lower expected future profitability for oil and energy producers. More interestingly, there is a strong positive effect on expected future profitability for a large section of the rest of the economy. In particular, expected future profits for producers of typical consumer goods industries typically increase when oil prices decrease. These findings have, potentially, also important implications for asset allocations.

JEL Classification: E44, G11, G12, Q02

*For the latest version, please see <http://goo.gl/JeBeuv>. This paper was previously circulated under the title “Portfolio Choice and Non-Tradable Commodity Risk”. We thank Borja Larrain, Thore Johnsen, Bernard Dumas, Pierre Collin-Dufresne, Kjetil Storesletten, and Ruey Tsay for their insightful discussions and comments. We also thank the participants at the Finance UC 11th International Conference, the 2016 European Meeting of the Econometric Society, and the seminar at the Norwegian School of Economics. The usual disclaimer applies.

[†]BI Norwegian Business School, email: espen.henriksen@bi.no

[‡]Norwegian School of Economics (NHH), email: jens.kvaerner@nhh.no

1 Introduction

Commodity price shocks may impact the level of aggregate economic activity and the profitability of different sectors of the economy. Competitive equity markets is a mechanism to converge on the present value of future profitability of publicly listed companies. The valuations of these companies may be aggregated at different levels, such as at the industry level or at the level of the aggregate economy.

Changes in commodity prices may affect the valuation of companies both through revisions of their future profitability and through the rate at which investors discount those future profits. In order to identify the information about changes in expected future industry profitability when equity-price and commodity prices move together it is crucial to separate the risk premium (or discount-rate) component and the profitability (or cash-flow) component. Changes in companies' discount rates change future expected returns. Changes to expected cash flows tend to be persistent shocks to the expected level.

Shiller (1981), Hansen and Singleton (1983), Mehra and Prescott (1985), Hansen and Jagannathan (1991), and many others showed how a wide class of models fail to simultaneously account for prices and quantities. Since then there has been a substantial progress in understanding the links between prices and quantities, in particular in accounting for prices given the dynamics of quantities (see e.g., Campbell and Cochrane, 1999, Bansal and Yaron, 2004, and Barro, 2006). This paper builds on these contributions and methodological innovations, but instead of accounting for prices given the dynamics of quantities, we strive to identify the information about future quantity dynamics given by changes in equity prices.

Evidence of time-varying discount rates goes back to at least Campbell (1987), Fama and French (1988), and others. Campbell and Shiller (1988) derive a loglinear approximate present-value relation that allows for time-varying discount rates. Campbell (1991) shows how to go from the log-linear present-value to decomposition of returns. Campbell and Mei (1993) break industry portfolio betas into components related to discount-rate news and cash-flow news. Bernanke and Kutter (2005) derive a procedure to assess how monetary policy surprises relate to expectations of future cash flows, future interest rates, and discount rate of the equity market. Building on these contributions, we develop a methodology to distinguish and quantitatively estimate how innovations to commodity prices relate to discount rate changes and future cash flows changes of different industries and sectors of the economy. We then apply this methodology to the price of crude oil, which is the largest commodity market, and the U.S. equity market, which is the world's largest equity market.

This paper also relates to previous research studying the relationship between oil price shocks and time-variation in expected returns. In a sample of 48 country indices, Driesprong et al. (2008) find that changes in oil prices predict aggregate market returns, and that the predictability is not spanned by variables well-known to predict returns, such as the term spread, dividend yield, default premium. Sørensen (2009) finds that only oil-price changes resulting from extreme events (such military conflicts in the Middle East, OPEC collapses, etc.) predict stock returns. These findings are broadly supported by Casassus and Higuera (2011), which finds that the stock market excess returns are significantly affected by oil price changes.

Conditional on the identification scheme, movements in commodity-prices may be decomposed into supply and demand shocks. Kilian (2009) identifies three components: supply shocks, aggregate demand shocks, and oil-specific demand shocks. Kilian and Park (2009) find that the response of aggregate stock returns may differ greatly depending on the cause of the oil price shock. Aastveit et al. (2014) find that demand from emerging economies is about twice as important as demand from developed countries in accounting for the fluctuations in the real oil price and in oil production. Ready (2014) suggests an alternative oil-price identification scheme relying only on traded variables and finds that the identified shocks to oil supply have larger impact on firms that depend on consumer expenditure than those that rely on oil as an input. That is broadly consistent with the findings of Lee and Ni (2002) and Gogineni (2010). We complement these papers by studying the relation between changes in current commodity prices and the long-term cross-sectional variation in profitability at the sector level as reflected by the asset markets.

Application: portfolio choice If changes in commodity prices covary with expected long-run profitability, either at the aggregate market level or at the industry level, that may have implications for financial portfolio choice, in particular for owners of *de facto* non-tradable commodity wealth. In a separate section, we analyze the asset allocation implications of our results.

This application follows in a long tradition in portfolio-choice theory. In the classic models of Mossin (1968), Samuelson (1969), and Merton (1969), with only tradable financial wealth and constant expected returns, the optimal equity versus bond split is independent of investors' time horizon. Implicitly, the within-equity-portfolio weights are time-invariant and equal to the market weights.

There is a large literature on portfolio choice with non-tradable assets. In particular, several papers have analyzed how portfolio compositions can partly offset the risk associated with

individuals' non-tradable current and future labor income. Bodie et al. (1992), Heaton and Lucas (1997), Viceira (2001), Storesletten et al. (2007), and Benzoni et al. (2007) incorporate non-tradable labor income into intertemporal models. These papers show that an individual's relative valuation of bonds and equities, and hence portfolio weights, change in response to the riskiness of her or his future labor income. A similar argument was made by van den Bremer et al. (2016) for sovereign wealth funds with respect to oil. In addition to analyzing the implications of non-tradable commodity income for the split between bonds and equities, in this paper we also analyze how it may affect the composition of the equity portion of the portfolio.

We find that for any owner of *de facto* non-tradable oil wealth it seems to be optimal to hold more equities and less bonds than it would have been if she/he only owned tradable financial assets. Further, it also seems optimal to hold less-than-market weights for the oil and energy sector, and higher-than-market weights in sectors such as production of nondurable consumer goods, and retail and wholesale. Interestingly, the potential gains from higher-than-market weights in production of nondurable consumer goods and in retail and wholesale, seem to be at least as large as the gains from holding less-than-market weights in the oil and energy sector.

2 Commodity Price Changes and Equity Returns

The present-value formula points to two reasons why equity prices may change: either expected cash flows change, discount rates change, or both. In this section, we derive a model to estimate how changes in commodity prices affects these two components of unexpected return for the aggregate equity market and for different industries. Returns are modelled as a vector autoregressive model with exogenous variables (VARX). Our approach shares some methodological similarities with Bernanke and Kutter (2005), which decompose the impact of monetary policy shocks on equity returns into a changes in discount rates, real interest rates, and expected cash-flows using an extended version of the return decomposition of Campbell (1991).

2.1 A Discount-Factor and Cash-Flow Component of Returns

Following Campbell and Shiller (1988), we express stock prices using a loglinear approximate present-value relation that allows for time-varying discount rates

$$P_t = E_t \sum_{j=1}^{\infty} \left(\prod_{k=1}^j \frac{D_{t+k}}{R_{t+k}} \right) \quad (1)$$

D is the dividend stream and R denotes the discount rate. To linearize Equation (1), we start with the trivial equality $1 = R_{t+1}R_{t+1}^{-1} = R_{t+1}^{-1} \frac{P_{t+1}+D_{t+1}}{P_t}$ and divide both sides by D_t to get

$$\frac{P_t}{D_t} = R_{t+1}^{-1} \left(1 + \frac{P_{t+1}}{D_{t+1}} \right) \frac{D_{t+1}}{P_t}$$

A first order Taylor expansion around the mean of the natural logarithm of the price-to-dividend ratio gives the return approximation

$$r_{t+1} \approx k + \Delta d_{t+1} + \rho(p_{t+1} - d_{t+1}) - (p_t - d_t) \quad (2)$$

where k is the constant from the log-linearization given by $k = \ln(1 + \frac{P}{D}) - \rho(p - d)$ and ρ is the slope of the log-linear function given by $\rho \approx 1 - \frac{D}{P}$. Iterating the return approximation forward, taking expectations, and imposing that the price-to-dividend ratio is bounded from above, we get the approximate identity

$$p_t - d_t \approx k_1 + E_t \sum_{j=1}^{\infty} \rho^{j-1} (\Delta d_{t+j} - r_{t+j}) \quad (3)$$

This equation is a result of the accounting identity relating stock prices to future expected cash flows discounted at the current discount rate. It holds ex post as well as ex ante. The log price-dividend ratio is high when investors expect high cash flow growth, or when stock returns are expected to be low. Campbell (1991) shows how to go from the log-linear present-value to decomposition of returns. Substituting Equation (3) into the approximate return equation gives:

$$\begin{aligned} r_t - E_{t-1}r_t &= (E_t - E_{t-1}) \left[\sum_{j=0}^{\infty} \rho^j \Delta d_{t+j} - \sum_{j=1}^{\infty} \rho^j \Delta r_{t+j} \right] \\ &= N_{CF,t} - N_{DR,t} \end{aligned} \quad (4)$$

N_{CF} is news about future cash flows (i.e., shock to expected earnings), and N_{DR} denotes news about future returns (i.e., shock to discount rate). This equation says that the unexpected part of the realized return, $(r_t - E_{t-1}r_t)$, stem from revision in expectations of future cash flows, changes in the discount rate, or both. The interpretation of the news components are the effects of today's shock over the discounted infinite future. Given the return decomposition above, the commodity-price exposure for equity portfolio j is

$$\beta_j = \frac{Cov(r_{jt} - E_{t-1}r_{jt}, N_{Ot})}{Var(N_{Ot})} = \frac{Cov(N_{CF,jt} - N_{DR,jt}, N_{Ot})}{Var(N_{Ot})} \equiv \beta_{CF}^O - \beta_{DR}^O, \quad (5)$$

where N_{Ot} is the time- t innovation in commodity price change. We term the regression slopes of discount-rate and cash-flow shocks on commodity price innovations as discount-rate exposure (β_{DR}) and cash-flow exposure (β_{CF}) to commodity prices.

2.2 Estimation Procedure

We estimate discount-rate exposure (β_{DR}) and cash-flow exposure (β_{CF}) to commodity prices with a VAR approach following Campbell et al. (2013); Campbell and Vuolteenaho (2004); Campbell and Mei (1993); Campbell (1991). The implicit assumption is that expected returns are linear in the state variables, which follow a first-order VAR.¹ In order to relate the proxies for expectations about future excess return ($N_{DR,t}$) and cash-flows ($N_{CF,t}$) to unanticipated changes in the commodity price (ΔZ_{t+1}), we estimate a vector autoregressive model with exogenous variables (VARX) model. After demeaning the variables, the VARX model is

$$\begin{aligned} \mathbf{X}_{t+1} &= \mathbf{A}\mathbf{X}_t + \mathbf{u}_{t+1} \\ &= \mathbf{A}\mathbf{X}_t + \mathbf{\Phi}\Delta Z_{t+1} + \mathbf{w}_{t+1}^\top, \end{aligned} \quad (6)$$

where \mathbf{X}_{t+1} is a $m \times 1$ vector of state variables with r_{t+1} as its first element, \mathbf{A} is the m -by- m estimated coefficient matrix, \mathbf{u}_{t+1} is an $m \times 1$ vector of residuals, and $\mathbf{\Phi}$ is the vector of surprise coefficients that capture the contemporaneous response of the elements of \mathbf{X}_{t+1} to an commodity price surprise at $t + 1$. To derive the discount-rate exposure (β_{DR}) and cash-flow exposure (β_{CF}) to commodity prices, let e_1 be a vector whose first element is equal to one and zero otherwise and I is the $m \times m$ identity matrix. Solving for discount rate news ($N_{DR,t}$) using

¹See Appendix A,

Equation (4) gives

$$\begin{aligned}
N_{DR,t+1} &= e1\lambda\mathbf{u}_{t+1} \\
&= e1\lambda(\Phi\Delta Z_{t+1} + \mathbf{w}_{t+1}^\top) \\
&= \underbrace{e1\lambda\Phi}_{\beta_{DR}^Q} \Delta Z_{t+1} + e1\lambda\mathbf{w}_{t+1}^\top
\end{aligned} \tag{7}$$

$$\begin{aligned}
N_{CF,t+1} &= (e1 + e1'\lambda)\mathbf{u}_{t+1} \\
&= (e1 + e1'\lambda)(\Phi\Delta Z_{t+1} + \mathbf{w}_{t+1}^\top) \\
&= \underbrace{(e1 + e1'\lambda)\Phi}_{\beta_{CF}^Q} \Delta Z_{t+1} + (e1 + e1'\lambda)\mathbf{w}_{t+1}^\top
\end{aligned} \tag{8}$$

where the function lambda $\lambda \equiv \rho\mathbf{A}(I - \rho A)^{-1}$ maps shocks to state variables to news about future cash flows and discount rate. The term $e1'\lambda\Phi$ can be interpreted as a vector of weights that determine the importance of commodity price shocks to discount rate expectations; the greater absolute value, the larger impact of a shock in the state variables.

We estimate the VARX model using a two-stage procedure. With this two-stage estimation, we can use the long time series of the equity market to estimate the coefficient matrix A , which will improve the estimation precision. The fact that commodity prices innovations from time t to $t + 1$ are uncorrelated with the level of the state variables at time t , ensures consistent estimates of both \mathbf{A} and Φ using two stage estimation.

2.3 Commodity Price Surprises and Equity Returns

3 Oil Price Surprises and U.S. Equity Returns

In the following, we apply our methodology to the oil market and the United States equity market. The reasons are primarily that the oil market is the world's largest and most important commodity market and that the U.S. equity market is the world's largest and most important equity market.

There is no *a priori* reason to assume that the covariance between equity valuations and oil prices should be the same across industries. This is also consistent with the findings of, among others, Lee and Ni (2002), Gogineni (2010), and Chiang et al. (2014). We therefore start with

the Fama and French 11 industry definitions (dropping “Other”). Together these industries constitute the market portfolio. Moreover, our reference oil prices is the crude oil futures front month (spot) contract from Chicago Mercantile Exchange and Chicago Board of Trade (CME) available for the period 1983:Q1-2015:Q4. The futures prices are unadjusted prices and reflect raw prices from CME.² All nominal oil prices has been deflated by the US Consumer Price Index for All Urban Consumers.³

We estimate a VAR model with three state variables at the quarterly horizon. The first variable in the VAR is the excess log return on the portfolio under scrutiny (r_j^e), that is, the difference between the log return on a value weighted portfolio described above and the log risk-free rate. The risk-free-rate data are obtained from CRSP, and represents the treasury bill closest to 90 days maturity in each month. The next state variable is the portfolio specific dividend-to-price ratio (DP). We create per-share dividends series from CRSP following the procedure in Campbell and Shiller (1988); Bansal et al. (2005); Bansal and Kiku (2011) among others. Specifically, let the total return per dollar invested be given by $R_{t+1} = R_{t+1}^X + Y_{t+1}$, where R_{t+1}^X is the capital accumulation and Y_{t+1} is the dividend yield. The level of dividends per share (D_{t+1}) and the share price (P_{t+1}) can then be computed as follows:

$$D_{t+1} = Y_{t+1}P_t \tag{9}$$

$$P_{t+1} = R_{t+1}^X P_t \tag{10}$$

with $P_0 = 1$. The interpretation of the dividend series that we use (D_{t+1}) is the total cash dividends received by an investor at $t + 1$ that extracts the dividends and reinvests the capital gains. For all the portfolios, the dividend series are constructed on a monthly basis and converted into quarterly resolutions by summing the level of dividends within a quarter. To reduce the impact on our estimates from the strong seasonality embedded in dividends, we follow the common practise in the literature and use a trailing four quarter average of the quarterly dividends as our final estimate of quarterly dividends. Finally, we construct the price-to-dividends ratio by dividing our smoothed dividends (D_{t+1}^S) on the the price level (P_{t+1}) the same quarter. The ratio is log transformed and is available from 1926-2014.

Our final state variable is obtained from taking the first principal component of most of the non-price quarterly predictors in Goyal and Welch (2008), and span the period 1952-2014. The first principal components account for about 30% of the variation in the data, and is normalized to have a mean of zero and a standard deviation of one. It is constructed from the following

²Data can be accessed through Quandl with code: CHRIS/CME_CL1

³Data can be accessed through Quandl with code: CPIAUCSL

variables: 1) The term spread (tms), computed as the difference between the long term yield on government bonds and the Treasury-bill. 2) The consumption–aggregate wealth ratio (cay) of Lettau and Ludvigson (2001). 3) Inflation defined as the Consumer Price Index (All Urban Consumers) from the Bureau of Labor Statistics. 4) The stock variance, computed as sum of squared daily returns on the S&P 500. 5) Investment–to–Capital ratio, defined as the ratio of aggregate (private nonresidential fixed) investment to aggregate capital for the whole economy. We refer the reader to the original article for a thorough description of the data.⁴

Table 1 reports summary statics and Figures 1-2 show the first principal component and the dividend yield and the following 3-year and 7-year log excess return, respectively.⁵ Both series use the value weighted market portfolio as the reference return. The figures indicate that both variable have some ability to predict future price levels at different horizons.

[Insert Table 1 here]

[Insert Figure 1 here]

[Insert Figure 2 here]

4 Results

4.1 Estimating Overall Covariances

First, we estimate the covariances between the oil price changes and the return on the aggregate equity market, and between the oil price changes and the return on separate industries, after controlling for their covariance with the aggregate market. Controlling for the covariance between oil price changes and the return on the aggregate market allows us to study the stability of industry specific oil price covariances over periods in which the oil and the equity market are driven by the same forces. Table 2 presents the results from regressing industry excess return on changes in oil price after controlling for the return on the aggregate market.

The first row in Table 2 shows that changes in oil prices are uncorrelated with the return on the market portfolio over the period 1983-2014 (similar results have been documented by Ready

⁴All data are available from Amit Goyal’s website: <http://www.hec.unil.ch/agoyal>.

⁵The first principal component series has been divided by 5 to ease comparison.

(2014), among others) but significantly positively correlated after the millennium with a coefficient of 0.13 (t -value of 3.7). After controlling for the market component, there are substantial heterogeneity among response coefficient across industries. In particular, the valuation of the energy sector is positively correlated with the oil price whereas the valuation of consumer goods industries and financials are negatively correlated with the oil price. For example, a 50% decrease in the oil price is associated with a 10% decrease in the market value of energy related equities (t -value between 7 and 11.8), after controlling for the market. The corresponding estimate for the the retail sector following a 50% decline in the oil price is about positive 5% (t -value 3.9-4.6). The estimated negative oil price exposure of typical consumer goods industries is broadly consistent with the hypothesis of Hamilton (2003) that oil price shocks act primarily on consumer expenditure, instead of pushing up input costs. The estimated slopes coefficients for industries with significant oil price exposure are stable across sample periods. Parameter stability across periods alleviate potential concern that a few extreme observations may have an unduly large impact on the results.

[Insert Table 2]

4.2 Oil Prices and Industry Profitability

Competitive equity prices reflect forecasts of future profits discounted at the equilibrium discount rate. Consequently, prices can move unexpectedly either because investors update their expectations of future cash flows, or change the risk premia they use to discount these cash flows, or both. Table 3 presents the results from estimation of the VAR for the aggregate market portfolio. Consistent with previous research (see e.g., Campbell, 1991) discount rate news is about twice as volatile as cash flow news for the market portfolio. Moreover, cash-flow and discount-rate news are weakly negatively correlated. The negative correlation between discount-rate and cash-flow news reflects that, on average, good news about future cash flows are accompanied by modest decline in discount rate.

Panel A reports VAR coefficient estimates, and Panels B and C describe how the first principal component used in the VAR is related to its component. As reported in Table 3, two out of our three state variables have predictive power in forecasting excess return at the quarterly horizon. The excess return has an insignificant autocorrelation coefficient of 0.064, indicating little degree of momentum. Moreover, a high dividend-to-price ratio (dp) predicts high future excess returns whereas the first principal component is negatively correlated with future excess return. The R^2 is about 6.0 and the F -test that all the predictors are jointly insignificant is rejected at the one percent level. The remaining rows in the VAR system show that the

predictive variables are fairly persistent processes with modest degree of co-dependence. Panel C reports the R^2 from regressing the first principal component on its component. The Panel shows that the term spread (tms) is the most influential factor whilst the stock variance (svar) is the least.

[Insert Table 3]

The portfolio-level VARs are presented in Table 4. The main results can be summarized as follows. The industry-specific VARs have all some predictive power in forecasting excess return at the quarterly horizon. Moreover, the lagged coefficients on both the industry specific dividend-to-price ratio (dp) and the first principal component have the same sign for all industries. The R^2 for the predictive regressions ranges from about 0.02–0.06 and the p -value of the F -test that all the predictors are jointly insignificant is rejected at the 10 percent significant level for all industries.

[Insert Table 4]

It is well-known that estimates of persistent AR(1) coefficients may be biased downwards in finite samples. The dividend-to-price ratio we use in the VAR is both persistent and their innovations are correlated with returns, which causes bias in the estimates of predictive regressions (Stambaugh, 1999). In an important endeavor of addressing whether the Stambaugh bias have material impact on the estimated discount-rate and cash-flow news Campbell et al. (2010) re-estimated their VAR system 2500 times. They find that the VAR coefficients are modestly biased, but that the bias have immaterial effect on their estimates of cash-flow and discount-rate news. We follow their approach and do not correct for the potential the bias.

Table 5 reports the results from estimating the discount-rate and the cash-flow exposure of commodity equity as well as for the market portfolio for the period 1983-2014, and the two subperiods 1983-1999, and 2000-2014.

[Insert Table 5]

The relation between changes in oil prices and future profitability is very stable across the sample periods with an estimate of negative -0.06 . The negative cash-flow component implies that increasing oil prices have a negative impact on long-run profitability in the United States equity market. As a result of the stable cash-flow component, the change in the sign of the covariance between changes in oil prices and the return on the market portfolio before and after

the millennium must stem from the discount factor component. Table 5 shows that for the period 1983-1999, increasing oil prices were associated with raising discount rates, whereas for the period after 2000, increasing oil prices were associated with declining discount rates. Figure 3 shows the price reaction for the US equity market following a 50% decline in the oil price as a function of time.

[Insert Figure 3 here]

The slowly increasing cumulative log return function reflects that declining oil prices over time are associated with higher profitability. The dashed line shows the same impulse response but estimated using oil price data from the period 2000-2015. The impulse response shows that the recent positive covariance between oil prices and the US equity market comes from the discount rate component. The more long-term consequences of changing oil prices for profitability have remained unchanged.

The estimates for the four industries reveal a similar pattern. The estimated relation between changes in oil prices and expected industry profitability is both industry specific and relatively stable over time. Not surprisingly, a negative oil-price shock is associated with lower future profitability for oil and energy producers. More interestingly, there is a strong positive effect for future profitability for producers of non-durable goods, retail and wholesale, and the financial industry. The estimated relation between changes in oil prices and industry discount rate is less industry specific and highly period dependent. Figure 4 shows the the price reaction for four industries following a 50% decline in the oil price as a function of time (energy sector is at the top left, the retail sector is at the top right, nondurable consumption goods industries are at the bottom left, and the financial sector is at the bottom right). The figure shows that profitability dynamics of changes in oil prices are stable over time and highly industry specific.

[Insert Fig 4 here]

5 Application to Asset Allocation

When returns are predictable, contemporaneous covariances are unsatisfactory risk measures for investment periods that exceeding one period. The reasons is that with stationary discount rates, increases in discount rates (risk premium) in the current period will be followed by higher expected returns in the future. In contrast, negative news about future cash flows are

permanent wealth loss independent of the investment horizon. Distinguishing between discount-rate and cash-flow risk is therefore essential for portfolio considerations, in particular for owners of de-facto non-tradable commodity wealth.

In this section, we use our estimates of discount-rate exposure and cash-flow exposure of equity portfolios to compute both optimal deviation from initial allocation between equity and bonds, and optimal deviations from market weights given that the market portfolio is mean-variance efficient. For the same reasons as mentioned above, oil wealth is an interesting example of a non-tradable asset for several reasons. First, it is arguably one of the worlds most important commodities and subject to substantial price fluctuations. Second, for regulatory reasons, the ownership is *a*o usually highly regulated and anchored to the country in which the oil extraction takes place. Together this makes oil important, oil reserves risky, and *de facto* non-tradable.⁶

5.1 The Setup

We focus on a marginal investor with non-tradable oil wealth who follows a buy-and-hold strategy over different holding periods. The impact of return predictability on portfolio choice for static buy-and-hold investors is very close to investors who myopically rebalance (Barberis, 2000). Our asset space consists of N tradable (TR) and 1 non-tradable asset (NT). The vector α and the scalar ω denote investor's positions in the tradable and non-tradable assets as fractions of initial financial wealth (FW). Total wealth (W_{TOT}) is therefore related to financial wealth (FW) and non-tradable wealth (NTW) through $W_{TOT} = (1 + \omega)FW$ where $\omega = NTW \times FW^{-1}$. The investor's portfolio constraint applies *de facto* only to the positions in tradable assets.

⁶Whether the composition of investors' equity-portfolios can improve the long-run risk-return ratio of total wealth for investors with substantial oil reserves may be in particular interest to oil-rich countries. Over the past decades, many oil-rich countries (Norway, the United Arab Emirates, etc.) have established large sovereign wealth funds (SWFs) from the oil income. The official objectives of many SWFs were to share the value of a non-renewable resource among an undefined number of future generations, and to provide diversification from commodity price risk. As of July 2014, natural resource funds hold approximately \$4.0 trillion in assets⁷, equivalent to about 30% of the US stock market. Empirical research shows that the SWFs tend to invest in large foreign firms, often in the finance and energy sectors, with low diversification and poor medium term performance (Bernstein et al., 2013; Chhaochharia and Laeven, 2009; Dyck and Morse, 2011; Bortolotti et al., 2015).

5.2 The Equity Share

In this context, at time t , the investor's asset demand is given by the following s -period portfolio problem

$$\max_{\alpha_{s,t}} \left\{ E[FW_{t+s}] - \frac{1}{2}\gamma Var[FW_{t+s}] \right\} \quad (11)$$

subject to

$$FW_{t+s} = \underbrace{FW_t[(1 + R_{f,t+s}) + \alpha' \mu_{TR}]}_{\text{Value of Tradable Assets}} + \underbrace{FW_t[\omega R_{nt,t+s}]}_{\text{Value of Non-Tradable Assets}}$$

Here, $\alpha_{s,t}$ is the vector of portfolio weights in tradable assets and γ is the coefficient of risk aversion. Denote $\mu_{TR} \equiv E[r_e]$ as the vector of excess return on tradable assets, $\Sigma_{TR} \equiv E[(r^e - \mu_{TR})(r^e - \mu_{TR})']$ as the variance-covariance-matrix of tradable assets, and let $\Sigma_{TR,NT}$ be a vector of covariances between returns on tradable and the non-tradable assets. In this setup, the investor's optimal portfolio weights are

$$\alpha = \gamma^{-1} \Sigma_{TR}^{-1} \mu_{TR} - \Sigma_{TR}^{-1} \Sigma_{TR,NT} \omega \quad (12)$$

The first term in Equation (12) represents the weights in the mean variance efficient portfolio, whilst the second component represents the hedging demand induced by investor's non-tradable asset. Because the financial market portfolio is mean variance efficient, Equation (12) collapses to standard two-fund separation plus an adjustment for non-tradable income

$$\alpha = \gamma^{-1} \frac{\mu_M^e}{\sigma_M^2} - \omega \frac{\sigma_O^2}{\sigma_M^2} \beta_{MO} \quad (13)$$

where $\beta_{MO} \equiv \frac{Cov(r_{Mt}^e, \Delta o_t)}{Var(\Delta o_t)}$ represents the regression coefficient from regressing log excess return of the market portfolio (r_{Mt}^e) on log changes in prices of the non-tradable asset (Δo_t). Substituting the decomposition in Equation (5) into Equation (13) gives

$$\alpha = \gamma^{-1} \frac{\mu_M^e}{\sigma_M^2} - \omega \frac{\sigma_O^2}{\sigma_M^2} (\beta_{MO,CF} - \beta_{MO,DR}) \quad (14)$$

Equation (14) shows that the investment horizon impacts the equity share directly. Long-term buy-and-hold investors ignoring temporary changes in prices should focus on the permanent component of the covariance ($\beta_{MO,CF}$). That is, how the value of the non-tradable asset covary with the expected profitability of the equity market. Figure 5 shows the risky share for an

investor with the same amount of wealth in non-tradable and financial wealth (i.e., $\omega = 1$) that would have an equity share of 50% in the absence of non-tradable assets using data from three different sample periods. The relative variance of oil and the equity market, which is calibrated to match the period 1983-2014, equals about 4.

[Insert Figure 5 here]

Investors with substantial *de facto* non tradable oil-revenue wealth should hold a larger proportion of equities than an investor with no non-tradable oil wealth. The reason is that oil wealth is negatively correlated with the long-run profitability of the economy. Interestingly, the portfolio suggested portfolio adjustment is remarkably stable over time. The largest difference between any of the three sample periods is only 0.03.

5.3 Composition of the Equity Portfolio

Yet the market portfolio is mean variance efficient, when different components of the market portfolio exhibit disparate sensitivity to fluctuations in commodity prices, investors with non-tradable commodity assets may improve the risk-return tradeoff of total wealth by deviating from market weights in the financial portfolio. Deviations depends on the investment horizon. If we for simplicity assume that the funds stemming from the underweighting are invested in the market portfolio, optimal deviation from the market weight reflects a trade off between two variances. The reduction in variance comes from removing a fraction of the industry with the highest correlation with the non-tradable asset from the financial portfolio whereas the additional variance comes from the efficiency loss caused by deviating from market weights, and the additional variance resulting from a potential non-zero covariance between the market portfolio and the non-tradable asset. Formally, it solves

$$\min_{\tilde{\alpha}} \left((\sigma^2(\tilde{\alpha}r_M + \omega\Delta O) + \sigma^2(-\tilde{\alpha}r_H + \omega\Delta O)) \right), 0 \tag{15}$$

with the solution to Equation (15) given by

$$\tilde{\alpha} = \psi\omega(\beta_{HO} - \beta_{MO}) \tag{16}$$

where $\beta_{jO} \equiv \frac{Cov(r_{jt}^e, \Delta op_t)}{Var(\Delta op_t)}$ represents the regression coefficient from regressing log excess return of portfolio j on changes in log prices of the non-tradable asset (Δop_t), $\psi \equiv \left(\frac{\sigma_O^2}{\sigma_M^2} \right) \left(\frac{R2}{1-R2} \right)$ is a constant that scales deviation from market weights according to the relative variance between

the non-tradable asset and the market portfolio, and the amount of idiosyncratic risk in the hedge portfolio. $R2$ is the coefficient of determination from a regression of the hedge-portfolio (H) on the market portfolio (M). Substituting the decomposition in Equation (5) into Equation (13) gives

$$\alpha = -\psi\omega (\beta_{HO,CF} - \beta_{MO,CF}) \quad (17)$$

Figure 6 illustrates the optimal deviations from market weights that comes out of Equation (17) for four different industries using data from the period 1983-2014 for an investor with the same amount of wealth in non-tradable and financial wealth (i.e., $\omega = 1$).

[Insert Figure 6 here]

Not surprisingly, investors with large oil assets show take out the energy sector from their portfolio. More interestingly, these investors should also invest considerably more in typical consumer goods industries. The retail sector is particularly attractive because of both a relatively high negative cash-flow exposure and little residual risk in the portfolio (i.e., a regression of the retail portfolio on the market portfolio gives an $R2$ above 0.7). The $R2$ is directly related to the amount of idiosyncratic risk in the hedge portfolio. The elasticity of the optimal deviation from market weight ($\tilde{\alpha}$) with respect to $R2$ is the fraction $\frac{\tilde{\alpha}}{(1-R2)}$. Using this elasticity, a reduction in the $R2$ of 20% (from 0,7 to 0,56 for the case of oil) reduces optimal deviation with approx. $1/1.2 - 1 \approx 17\%$. Similarly, an increase of 20% will increase the optimal deviation with $1/0.8 - 1 \approx 25\%$.

6 Conclusion

Changes in commodity prices may be related to both changes in risk premia and to the long-run aggregate level of economic activity and long-run profitability of different industries. Since at least Hansen and Singleton (1983), asset pricing researchers have strived to simultaneously account for prices and quantities. The majority of the research since then have attempted to account for prices given the dynamics of quantities. In this paper we build on some of these contributions to identify information about future quantities given price changes. In particular, using equity markets data and modern asset pricing methodology we separate two effects of changes in commodity prices: the long-run effect on profitability and the immediate effect on risk premia and how these profits are discounted.

Applying this methodology to the oil market, which is the world's largest commodity market, and the U.S. equity market, which is the world's largest equity market, we find that most of the short-run equity price movements associated with oil-price news can be accounted for by discount-factor effects. The discount-factor effects are, however, unstable. The long-run cash-flow, or profitability, effects are in general both more stable and statistically significant. A positive shock to oil prices are associated with lower overall expected profitability of the aggregate of listed companies.

Analyzing the effect on industry level, we find, again, that most of the immediate movements in equity prices following an unexpected oil-price change are due to changes in risk premia, ie. in the discount factor. The long-run profitability, or cash flow, effects are more stable and statistically significant. Non surprisingly, the expected profitability of the oil and energy industry is positively related to non-expected increases in the oil price. Perhaps more interestingly, the profitability of producers of non-durable goods, of retail and wholesale, and of the financial industry seem all to be negatively related to oil-price increases.

In addition to being a contribution to a deeper understanding of the workings of the macroeconomy, these findings may also be important for financial portfolio considerations. For an investor with substantial *de facto* non-tradable commodity wealth, these results have implications for composition of the financial tradable portfolio, both for the overall equity-bond allocation and for the composition of the equity portion of the portfolio.

Applying the financial portfolio methodology to the oil market and the US financial markets, an investor with substantial *de facto* non-tradable oil-revenue wealth may consider holding a larger proportion of equities than an investor with no non-tradable oil wealth. The reason is that oil wealth is negatively correlated with the expected long-run profitability of the economy. This investor should also, almost trivially, decrease her or his exposure to the oil and energy industry. More interestingly, her or his potential gains from increasing the exposure to producers of non-durables, retail and wholesale, and the financial industry are larger than the gains from reducing the exposure to oil and energy.

References

- Aastveit, K. A., Bjørnland, H. C., and Thorsrud, L. A. (2014). What drives oil prices? Emerging versus developed economies. *Journal of Applied Econometrics*, pages 1099–1255.
- Bansal, R., Dittmar, R. F., and Lundblad, C. T. (2005). Consumption, dividends, and the cross-section of equity returns. *Journal of Finance*, 60(4):1639–1672.
- Bansal, R. and Kiku, D. (2011). Cointegration and long-run asset allocation. *Journal of Business & Economic Statistics*, 29(1).
- Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance*, 55:225–264.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866.
- Benzoni, L., Collin-Dufresne, P., and Goldstein, R. S. (2007). Portfolio choice over the life-cycle when the stock and labor markets are cointegrated. *Journal of Finance*, 62(5):2123–2167.
- Bernanke, B. S. and Kutter, K. N. (2005). What explains the stock market’s reaction to Federal Reserve policy? *Journal of Finance*, 60(3):1221–1257.
- Bernstein, S., Lerner, J., and Schoar, A. (2013). The investment strategies of sovereign wealth funds. *Journal of Economic Perspectives*, 27(2):219–238.
- Bodie, Z., Merton, R. C., and Samuelson, W. F. (1992). Labor supply flexibility and portfolio choice in a life-cycle model. *Journal of Economic Dynamics and Control*, 16:427–449.
- Bortolotti, B., Fotak, V., and Megginson, W. L. (2015). The sovereign wealth fund discount: Evidence from public equity investments. *Review of Financial Studies*.
- Campbell, J., Giglio, S., and Polk, C. (2013). Hard times. *Review of Asset Pricing Studies*, 3(1):95–132.
- Campbell, J. Y. (1987). Stock returns and the term structure. *Journal of Financial Economics*, 18(2):373–399.
- Campbell, J. Y. (1991). A variance decomposition for stock returns. *Economic Journal*, 101(405):157–179.

- Campbell, J. Y. and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107:205–251.
- Campbell, J. Y. and Mei, J. (1993). Where do betas come from? Asset price dynamics and the sources of systematic risk. *Review of Financial Studies*, 6(3):567–592.
- Campbell, J. Y., Polk, C., and Vuolteenaho, T. (2010). Growth or glamour? Fundamentals and systematic risk in stock returns. *Review of Financial Studies*, 23(1):305–344.
- Campbell, J. Y. and Shiller, R. J. (1988). Stock prices, earnings, and expected dividends. *Journal of Finance*, 43(3):661–676.
- Campbell, J. Y. and Vuolteenaho, T. (2004). Bad beta, good beta. *American Economic Review*, 94(5):1249–1275.
- Casassus, J. and Higuera, F. (2011). Stock Return Predictability and Oil Prices. Documentos de Trabajo 406, Pontificia Universidad Católica de Chile.
- Chhaochharia, V. and Laeven, L. (2009). The investment allocation of sovereign wealth funds. Mimeo, University of Miami.
- Chiang, I.-H. E., Hughen, W. K., and Sagi, J. S. (2014). Estimating oil risk factors using information from equity and derivatives markets. *Journal of Finance*.
- Driesprong, G., Jacobsen, B., and Maat, B. (2008). Striking oil: Another puzzle? *Journal of Financial Economics*.
- Dyck, A. I. J. and Morse, A. (2011). Sovereign wealth fund portfolios. Research Paper 11-15, Chicago Booth.
- Fama, E. F. and French, K. R. (1988). Permanent and temporary components of stock prices. *Journal of Political Economy*, 96(2):246–273.
- Gogineni, S. (2010). Oil and the stock market: An industry level analysis. *The Financial Review*, 45(4):995–1010.
- Goyal, A. and Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21(4):1455–1508.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2):369–398.
- Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–62.

- Hansen, L. P. and Singleton, K. J. (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91(2):249–65.
- Heaton, J. and Lucas, D. (1997). Market frictions, savings behavior, and portfolio choice. *Macroeconomic Dynamics*, 1:76–101.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–69.
- Kilian, L. and Park, C. (2009). The impact of oil price shocks on the U.S. stock market. *International Economic Review*, 50(4):1267–1287.
- Lee, K. and Ni, S. (2002). On the dynamic effects of oil price shocks: a study using industry level data. *Journal of Monetary Economics*, 49(4):823–852.
- Lettau, M. and Ludvigson, S. (2001). Consumption, aggregate wealth, and expected stock returns. *Journal of Finance*, 56(3):815–849.
- Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.
- Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics*, 51:247–257.
- Mossin, J. (1968). Optimal multiperiod portfolio policies. *Journal of Business*, 41:215–229.
- Ready, R. C. (2014). Oil prices and the stock market. Working paper, University of Rochester.
- Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51:239–246.
- Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–36.
- Sørensen, L. Q. (2009). Oil price shocks and stock return predictability. Working paper, NHH.
- Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics*, 54(3):375–421.
- Storesletten, K., Telmer, C. I., and Yaron, A. (2007). Asset pricing with idiosyncratic risk and overlapping generations. *Review of Economic Dynamics*, 10(4).

van den Bremer, T., van der Ploeg, F., and Wills, S. (2016). The elephant in the ground: Managing oil and sovereign wealth. *European Economic Review*, 82:113 – 131.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance*, 56(2):433–470.

A Appendix

A.1 Proof

Substituting (2) into the approximate return equation gives:

$$\begin{aligned}
 r_t - E_{t-1}r_t &= (E_t - E_{t-1})\left[\sum_{j=0}^{\infty} \rho^j \Delta d_{t+j} - \sum_{j=1}^{\infty} \rho^j \Delta r_{t+j}\right] \\
 &= N_{CF,t} - N_{DR,t}
 \end{aligned} \tag{18}$$

Equation (4) shows that unexpected returns ($r_t - E_{t-1}r_t$) are linear in (discounted) revision in expectation ($E_t - E_{t-1}$) of cash-flow growth and future excess returns. If we can predict future excess return, we can back out the cash-flow component as the residual (since we observe unexpected returns ($r_t - E_{t-1}r_t$)). Since the conditional expectation of future returns ($E_{t-1}r_t$) is linear in the predictive variables, we can use a VAR:

$$\mathbf{X}_{t+1} = \mathbf{A}_0 + \mathbf{A}\mathbf{X}_t + \mathbf{u}_{t+1} \tag{19}$$

Now, we solve for discount rate news ($N_{DR_{t+1}}$):

$$\begin{aligned}
 -N_{DR_{t+1}} &= -(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j \Delta r_{t+j} \\
 &= -E_{t+1} \sum_{j=1}^{\infty} \rho^j \Delta r_{t+j+1} + E_t \sum_{j=1}^{\infty} \rho^j \Delta r_{t+j+1} \\
 &= -e1' \sum_{j=1}^{\infty} \rho^j A^j X_{t+1} + e1' \sum_{j=1}^{\infty} \rho^j A^{j+1} X_t \\
 &= -e1' \sum_{j=1}^{\infty} \rho^j A^j (AX_t + \mathbf{u}_{t+1}) + e1' \sum_{j=1}^{\infty} \rho^j A^{j+1} X_t \\
 &= -e1' \rho A (I - \rho A)^{-1} \mathbf{u}_{t+1} \\
 &\equiv -e1' \lambda \mathbf{u}_{t+1}
 \end{aligned} \tag{20}$$

The function λ is the function that maps shocks to returns and state variables (\mathbf{u}_{t+1}) to discount rate news and cash flow news.

A.2 Tables

Table 1: Summary Statistics

A. Log Excess Returns					
Industry	Mean	Median	Std.	Min	Max
US Equity Market	0.006	0.012	0.045	-0.264	0.118
Nondurable	0.008	0.009	0.041	-0.242	0.133
Durable	0.004	0.007	0.069	-0.396	0.355
Manufacturing	0.007	0.014	0.055	-0.343	0.191
Energy	0.006	0.007	0.053	-0.208	0.173
Chemicals	0.007	0.011	0.046	-0.288	0.134
Business and Equipment	0.004	0.010	0.071	-0.306	0.184
Telecom	0.005	0.012	0.051	-0.178	0.192
Utilities	0.006	0.010	0.039	-0.135	0.106
Wholesale and Retail	0.007	0.010	0.051	-0.338	0.120
Health	0.007	0.010	0.046	-0.235	0.149
Financial Sector	0.006	0.012	0.056	-0.254	0.157
B. Predictors					
dp_M	-4.697	-4.648	0.411	-5.740	-3.640
tms	0.017	0.016	0.014	-0.035	0.045
cay	-0.000	-0.003	0.022	-0.051	0.040
$infl$	0.009	0.007	0.008	-0.034	0.042
$svar$	0.006	0.003	0.010	0.000	0.113
ik	0.036	0.035	0.004	0.028	0.044
C. Log Changes in Real Oil Prices					
ΔOP_{NYMEX}	-0.01	0.01	0.20	-0.92	0.82

The table reports the summary statistics of the excess log return for Fama-French 11 industry portfolios over the sample period 1983:Q1-2014:Q4, the natural logarithm of the dividend-to-price ratio of the market portfolio, the components of the first principal component for the period 1952:Q1-2014:Q4, and the first differences in real log oil prices for the period 1983:Q2 (we use the crude oil futures front month contract from CME as our reference oil price). The Term Spread (tms) is computed as the difference between the long term yield on government bonds and the Treasury-bill. The consumption- aggregate wealth ratio (cay) is from Lettau and Ludvigson. Inflation ($infl$) is defined as the Consumer Price Index (All Urban Consumers) from the Bureau of Labor Statistics. The Stock Variance ($svar$) is computed as sum of squared daily returns on the S&P 500. Investment to Capital Ratio (ik) is defined as the ratio of aggregate investment to aggregate capital for the whole economy.

Table 2: Industry Portfolios and Oil Price Exposure

$$r_{it}^e = \alpha + \gamma_i \Delta o p_t + \beta_i r_{mt}^e + u_{it}$$

Sample Period	1983-2014		1983-1999		2000-2014	
Industry (<i>i</i>)	γ	$t(\gamma)$	γ	$t(\gamma)$	γ	$t(\gamma)$
US Equity Market ($\beta = 0$)	0.03	1.24	-0.06	-1.81	0.14	3.86
Nondurable	-0.06	-4.28	-0.08	-4.44	-0.00	-0.15
Durable	-0.02	-0.76	-0.06	-2.46	-0.00	-0.01
Manufacturing	0.02	1.47	-0.02	-1.66	0.05	2.61
Energy	0.23	11.83	0.23	9.95	0.24	7.04
Chemicals	-0.03	-1.76	-0.03	-1.64	0.01	0.34
Business	0.02	0.89	0.02	0.83	-0.01	-0.37
Telecom	-0.00	-0.13	0.03	1.42	-0.07	-2.67
Utilities	0.00	0.19	-0.03	-1.27	0.04	1.12
Wholesale	-0.09	-6.51	-0.07	-3.89	-0.09	-4.16
Health	-0.05	-2.92	-0.03	-1.40	-0.03	-1.24
Financial Sector	-0.06	-3.91	-0.06	-3.43	-0.07	-2.36

The table reports the results from regressing the excess log return of Fama-French 11 industry portfolios on the excess log return of market portfolio and changes in the natural logarithm of real oil prices over different sample periods. For each sample period, the first column (γ) reports the industry coefficient on changes in the natural logarithm of real oil prices and the second column ($t(\gamma)$) reports the corresponding t-statistics. The first row, "US Equity Market", reports the coefficient for the market portfolio on changes in the natural logarithm of real oil prices.

Table 3: VAR Estimation Results, Market Portfolio

	1	2	3	R2	F-Test	$e1\lambda$
Market Portfolio	Coefficients Matrix					
A. Slopes (t-statistics in parentheses p-values in brackets)						
Log excess return	0.064 (1.0)	0.026 (2.1)	-0.017 (3.3)	0.06	5.2 [0.00]	0.04
Log Dividend-to-Price ratio	-0.043 (0.7)	0.972 (76.8)	0.015 (2.8)	0.96	2 028 [0.00]	0.73
First Principal Component	-0.014 (0.0)	-0.016 (0.2)	0.882 (29.0)	0.78	287 [0.00]	-0.05
B. R2 From Regressing the First Principal Component onto its Components						
R2		tms	cay	infl	svar	ik
		0.61	0.26	0.48	0.08	0.56
C. Correlation Matrix						
Correlations	PCA	tms	cay	infl	svar	ik
PCA	1.00	-0.78	-0.51	0.69	-0.28	0.75
tms	-0.78	1.00	0.25	-0.33	0.15	-0.49
cay	-0.51	0.25	1.00	-0.21	0.01	-0.21
infl	0.69	-0.33	-0.21	1.00	-0.21	0.35
svar	-0.28	0.15	0.01	-0.21	1.00	-0.02
ik	0.75	-0.49	-0.21	0.35	-0.02	1.00

Panel A shows the results obtained with a first-order VAR model including the log excess return, the first principal component obtained from a handful of the quarterly predictors in Goyal and Welch (2008), and the log of the market dividend-to-price ratio (dp). All variables have been demeaned. The data span from 1952:Q1 to 2014:Q4, resulting in 252 quarterly observations. The upper panel reports the estimated VAR coefficient matrix, the corresponding t-statistics, the R2 of each regression, and the F-test from testing whether all the predictors are zero for each variable in the VAR. T-statistics are based on unadjusted OLS standard errors. The last column on right hand side shows the function $e1\lambda$ that maps shocks to state variables to news about future excess returns. In this function, $e1$ is a vector with the first element equal to unity and the remaining elements equal to zero and $\lambda \equiv \rho A(I - \rho A)^{-1}$ and $\rho = 0.96^{1/4}$, where A is the point estimate of the VAR transition matrix, and ρ is the slope of the log-linearization coefficient, which we set to 0.96 per annum. Panel B reports the R2 from regressing the first principal component on its component while Panel C shows the correlation matrix for the principal components.

Table 4: VAR Estimation Results, Industry Portfolios

	1	2	3	R2	F-Test
<hr/>					
Energy Portfolio	Coeff. on Lagged Variable				
Slopes (t-statistics in parentheses) and [p-values in brackets]					
Log Excess Return	0.076 (1.2)	0.025 (1.7)	-0.011 (2.0)	0.032	2.8 [0.04]
Log Dividend-to-Price ratio	-0.984 (39.0)	0.986 (164.0)	0.001 (0.4)	0.991	9 265
First Principal Component	0.190 (0.6)	0.033 (0.4)	0.881 (29.4)	0.778	288
<hr/>					
Nondurable Portfolio	Coeff. on Lagged Variable				
Slopes (t-statistics in parentheses) and [p-values in brackets]					
Log Excess Return	-0.037 (0.6)	0.042 (2.6)	-0.016 (3.0)	0.05	4.7 [0.00]
Log Dividend-to-Price ratio	-0.974 (30.5)	0.988 (120.9)	-0.003 (1.2)	0.98	4 992
First Principal Component	-0.521 (1.4)	-0.051 (0.5)	0.875 (28.8)	0.78	291
<hr/>					
Retail Portfolio	Coeff. on Lagged Variable				
Slopes (t-statistics in parentheses) and [p-values in brackets]					
Log Excess Return	-0.032 (0.5)	0.020 (2.0)	-0.020 (3.2)	0.051	4.4 [0.01]
Log Dividend-to-Price ratio	-0.989 (24.1)	0.994 (152.5)	-0.004 (1.0)	0.990	7 909
First Principal Component	-0.133 (0.4)	-0.005 (0.1)	0.879 (28.7)	0.777	287
<hr/>					
Financial Sector Portfolio	Coeff. on Lagged Variable				
Slopes (t-statistics in parentheses) and [p-values in brackets]					
Log Excess Return	0.055 (0.9)	0.027 (1.7)	-0.013 (2.0)	0.03	2.5 [0.06]
Log Dividend-to-Price ratio	-0.979 (29.6)	0.979 (117.5)	0.008 (2.3)	0.98	4 890
First Principal Component	-0.245 (0.8)	-0.068 (0.9)	0.883 (29.1)	0.78	289

The table shows the results obtained with a first-order VAR model including the log excess return, the first principal component obtained from a handful of the quarterly predictors in Goyal and Welch (2008), and the log of the industry dividend-to-price ratio (dp). All variables have been demeaned. The data span from 1952:Q1 to 2014:Q4, resulting in 252 quarterly observations. The upper panel reports the estimated VAR coefficient matrix, the corresponding t-statistics, the R2 of each regression, and the F-test from testing whether all the null hypothesis that all the predictors are zero for each variable in the VAR. T-statistics are based on unadjusted OLS standard errors. The industry portfolios are selected from Fama-French 11 industry portfolios.

Table 5: Decomposing Industry Exposure to Oil Prices (Quarterly)

Portfolio	Market			Energy			NonDurable			Retail			Financial Sector		
	I	II	III	I	II	III	I	II	III	I	II	III	I	II	III
Slopes (OLS standard errors in parentheses) and [bootstrap standard errors in brackets]															
$b_{CF} - b_{DR}$	-0.01	-0.13	0.16	0.23	0.12	0.38	-0.08	-0.19	0.09	-0.11	-0.21	0.05	-0.05	-0.18	0.14
$N_{CF,t} = a + \beta_{CF} \Delta OP_t + u_t$															
β_{CF}	-0.06	-0.06	-0.06	0.07	0.03	0.13	-0.07	-0.05	-0.09	-0.13	-0.12	-0.14	-0.06	-0.11	0.03
$se(\beta_{CF})$	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)	(0.04)	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)	(0.05)	(0.03)	(0.03)	(0.05)
$se(\beta_{CF})$	[0.01]	[0.01]	[0.03]	[0.02]	[0.02]	[0.05]	[0.03]	[0.04]	[0.04]	[0.03]	[0.04]	[0.05]	[0.03]	[0.03]	[0.05]
$N_{DR,t} = a + \beta_{DR} \Delta OP_t + u_t$															
β_{DR}	-0.04	0.07	-0.22	-0.16	-0.10	-0.25	0.01	0.13	-0.18	-0.02	0.09	-0.19	-0.01	0.07	-0.12
$se(\beta_{DR})$	(0.03)	(0.04)	(0.05)	(0.02)	(0.03)	(0.03)	(0.03)	(0.04)	(0.05)	(0.04)	(0.05)	(0.05)	(0.03)	(0.03)	(0.05)
$se(\beta_{DR})$	[0.04]	[0.03]	[0.05]	[0.02]	[0.03]	[0.03]	[0.04]	[0.04]	[0.04]	[0.04]	[0.05]	[0.05]	[0.03]	[0.03]	[0.05]

The table reports the impact of log changes in real oil prices (ΔOP_t) on the unexpected excess return of industry portfolios and the market portfolio. Time period I refers to the full sample, while II refers to 1983:Q2-2014:Q4 and 1974:Q2-2014:Q4 for NYMEX and EIA prices, respectively, and III refers to the period 2000:Q1-2014:Q4.

The three-variable VAR(1) used to construct the "Discount Rate News" (N_{DR}) and the "Cash Flow News" (N_{CF}) constitute of the log excess returns, the portfolio dividend-to-price ratio (DP), and the first principal component obtained from a handful of the quarterly predictors in Goyal and Welch (2008). Further, the dependent variables are outcome from the VARX model: $\mathbf{X}_{t+1} = \mathbf{A}_0 + \mathbf{A}_1 \mathbf{X}_t + \Phi \Delta OP_{t+1} + \mathbf{w}_{t+1}^\top$, where Φ is the vector of surprise coefficients that capture the contemporaneous response of the elements of \mathbf{X}_{t+1} to an oil price surprise at $t+1$. The impact of 1 percent in change in the oil price on the discounted sums of expected future excess returns is $N_{DR,t+1} = e1\lambda u_{t+1} = e1\lambda(\Phi \Delta OP_{t+1} + \mathbf{w}_{t+1}^\top = e1\lambda\Phi$. Similarly, the implied response of the present value of current and expected future cash flows is $(e1 + e1\lambda)\Phi$. Standard errors are estimated using OLS and with bootstrap using 2000 replications with replacement.

A.3 Figures




Figure 1: The First Principal Component and Following 3-Year Return




Figure 2: Dividend yield and Following 7-year Return



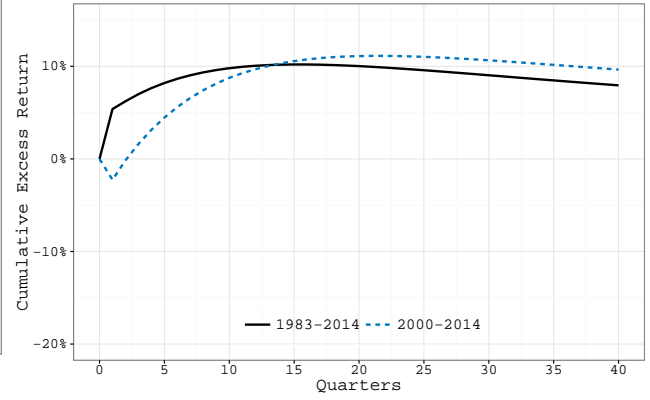



Figure 3: Market Portfolio




Energy Sector



Retail Sector



Nondurable Consumption



Financial Sector

Figure 4: Industry Portfolios




Figure 5: Allocation: Equity Share




Figure 6: Allocation: Sectoral Composition of an Equity Portfolio