Fracking, Drilling, and Asset Pricing: Estimating the Economic Benefits of the Shale Revolution

Erik Gilje* Robert Ready† Nikolai Roussanov‡

May 12, 2015

Abstract

We use evidence from asset price data to quantify the contribution of shale oil to the U.S. economy. Equity market valuations of firms engaged in shale oil extraction reflect the market's expectations about the future growth in shale oil supply and its potential for raising aggregate productivity in the U.S. economy. We pursue two complementary methods to estimate the value from shale oil discoveries. First, we examine returns on an index of shale oil producers orthogonalized with respect to oil prices and industry-wide return controls to extract an empirical measure of shale-specific productivity innovations. Second, we use the cross-section of stock returns that captures the variation in industries' responses to shale-specific news announcements to construct a shale-factor mimicking portfolio. While the two methods produce a wide range of estimates, overall they suggest that approximately 20%-25% of the increase in aggregate U.S. equity market capitalization since 2009 can be attributed to shale oil, corresponding to $2.5 to $3 trillion in value.

Keywords: cash-flow news, long-run growth, oil prices, shale oil, fracking, horizontal drilling

JEL codes: G12, G13, Q43

*The Wharton School, University of Pennsylvania

†Simon School of Business, University of Rochester

‡The Wharton School, University of Pennsylvania, and NBER

1 Introduction

Asset pricing theory is typically agnostic about the nature of technology shocks that underpin the variation in asset values.[1] At the same time, much of the debate in empirical research centers on the relative role of news about future cash flows in explaining variation in aggregate asset prices, as opposed to news about discount rates (e.g., Bansal and Yaron (2004), Campbell and Vuolteenaho (2004), Hansen, Heaton and Li (2008), Cochrane (2011), Greenwald, Lettau and Ludvigson (2014)). Over the five years following the Great Recession (2009 through 2014) the U.S. equity market capitalization roughly doubled, despite fairly anemic rates of growth in the real economy, suggesting falling discount rates as the main driver of rising valuations. However, over the same time period U.S. oil production increased dramatically, from less than 5 Mb/d (million of barrels per day) in 2010 to over 8 Mb/d in 2014, with total U.S. oil production forecast to nearly double by 2015 relative to the pre-crisis levels. Almost all of this increase can be attributed to a breakthrough technological innovation that allow oil to be extracted from shale rock formations that were previously thought to be too costly to access. This innovation, which involves the combination of hydraulic fracturing ("fracking") and horizontal drilling, in the matter of a few years has fundamentally changed the global energy supply-demand balance. Its success was also largely unexpected, as evidenced by the published forecasts of the Energy Information agency (EIA). Given the importance of oil to the U.S. economy, how much of the recent rise in the equity market can be attributed to the unexpected development of U.S. shale oil? Might this suggest a greater role for cash-flow news in explaining asset price fluctuations?[2]

Our approach centers on constructing asset market-based proxies for changes in expectations of shale oil productivity using both the time series and cross-section of industry returns. We pursue two complementary methods.

In the first approach we use stock returns of companies concentrated in shale oil extraction in order to construct the Shale Oil Index. We use this Shale Index together with the aggregate

[1] Recent work by Kogan, Papanikolaou, Seru and Stoffman (2012) linking news on patented technologies to equity returns is a notable and important exception.

[2] Our work here also fits into a long literature attempting to quantify the economic impact of oil shocks. Examples include Hamilton (1983), Sadorsky (1999), Hamilton (2003), Barsky and Kilian (2004), Kilian (2009), Kilian and Park (2009), Bodenstein, Guerrieri and Kilian (2012), and numerous others

market portfolio returns to estimate the contribution of the technological innovations in the shale sector to the broad equity market. We construct our index by including firms that have significant exposure to shale oil discoveries, as indicated by asset holdings listed in the firms' 10-K filings. Given the focus that these firms have on shale oil, changes in the index value should reflect the market's assessment of shale oil development.

The main challenge is to identify the component of stock returns that is attributable to technology shocks in the shale oil sector. There are several sources of shocks that contribute to both the aggregate stock market returns and the Shale Oil Index returns. Aggregate shocks to productivity in sectors that rely on petroleum as an input or increase household demand for oil through complementarities or income effects are generally good news both for the aggregate stock market and for oil prices, and, therefore, for oil producers. This can confound the effect of innovation in oil production. Shocks to other oil producers - domestic and foreign - as well as demand shocks abroad - have an opposite impact on the shale producers relative to the aggregate economy. Decreasing (net) supply to the U.S. market has a negative effect on the aggregate market but is a boon to domestic producers that are not themselves affected by the shock directly (a prominent example is the explosion on the Deepwater Horizon offshore rig and the subsequent oil spill in the Gulf of Mexico, which affected supply by BP, but was also potentially bad news for other offshore oil producers in light of the heightened safety concerns).

We address this challenge by orthogonalizing the Shale Oil Index returns with respect to both oil prices and the returns to an index of non-shale oil producers (North-American Integrated Index). We also orthogonalize the Shale Index return with respect to an index of shale gas producers, given the fact that some of the shale oil producers are also active in natural gas extraction and that could bias our inference about the contribution of the oil-extraction technology. The unexplained component representing innovations unique to shale oil firms we refer to as the "Shale shock". Finally, we estimate covariances between the Shale shock and the aggregate stock market portfolio. We find that roughly 20% of the increase in aggregate stock market returns since 2009 can be explained by our Shale shock, suggesting a gain of roughly $2.5 trillion based on the stock market appreciation over the period.

Increased shale production has two potentially positive effects for the U.S. Economy, one is the direct effect of increased employment and production in the Shale Oil sector, and the second is the indirect effect on oil prices, potentially lowering both their levels and long-run uncertainty about oil. Since we control for prices in our orthogonalization-based identification strategy, we are likely underestimating the total benefit from shale oil. The approach also suffers from potential endogeneity, as it is difficult to identify the precise source of correlation between the Shale Index and aggregate market returns (e.g., it could be that the Shale shock is picking up a greater sensitivity of shale stock returns to changes in discount rates).

Our second approach exploits the behavior of stock returns on days with major news announcements pertaining to the shale oil industry, as they are exogenous events with respect to aggregate shocks originating elsewhere in the macroeconomy. We first construct a series of exogenous shocks to shale expectations from earnings announcements of major shale producers. We find that unexpected positive news for shale producers leads to significant abnormal returns for shale firms, which have a significant positive impact on the aggregate market return. While this instrumental variable strategy enables us to cleanly identify the impact of shale production shocks on the market as a whole, the small number of these shocks does not allow for a direct estimate of the total benefit from shale oil, as these shocks presumably represent a fraction of the total change in market expectations over the period.

In order to address this issue, we examine the cross-section of industry returns (excluding the shale index). This allows us to exploit the cross-sectional variation in industries' exposure to shale oil news shocks by estimating their abnormal returns on two of the largest shale news announcements days. We find that exposure to shale news announcements has strong explanatory power for the cross-section of industry returns over the 2009 - 2014 period in which market exposure to the shale index is high, with industries that do well on the announcements of major shale discoveries earning a large premium over the period. This pattern is not there in the pre-crisis sample covering 2003 to 08/2008, and in fact we see a slightly negative relation between shale exposure and returns over this period.

Building on these results, we construct a factor-mimicking portfolio from the cross-section of industry portfolios that has maximum exposure to shale news shocks - the "Shale Factor." We use weights based on the shale announcement day returns and returns on the day of a key OPEC announcement in November 2014 that drive down oil prices significantly. This approach allows us to isolate industries that are positively (or negatively) exposed to shale oil news for reasons other than oil demand, since we are controlling for the response to the oil price itself with the OPEC announcement.

This Shale Factor-mimicking portfolio exhibits behavior which is consistent with results we find directly using the Shale Index. This portfolio is weakly correlated with the total stock market but negative returns over the pre-crisis sample period. After the crisis, the portfolio becomes increasingly correlated with the aggregate market while simultaneously earning extremely high returns (in-sample Sharpe Ratio ≈ 0.81) after 2009. The returns to the portfolio combined with the increased exposure of the aggregate market means that it can explain 25% of the market return since the financial crisis (i.e., implying a total contribution of about \$3 trillion). The high returns on the Shale Factor do not appear to be simply a product of increased market exposure, as the aggregate market return can only explain 21% of the returns to the Shale Factor itself. These results are somewhat l in line with those obtained using the Shale Shock. Thus, both of our main empirical strategies indicate that there is a meaningful economic impact of shale development.

Our estimates are likely to represent the lower bound on the contribution of the shale revolution on the U.S. economy. The empirical strategies that we pursue do not attempt to measure the impact of shale through the oil price itself, as both strategies are explicitly orthogonalized to oil prices.[3] How might the magnitudes we have found compare to the magnitude of the price effect? To put this comparison in context we undertake a simple back-of-the-envelope calculation. Total U.S. consumption of crude oil and petroleum products is approximately 18 Mb/d. Assuming that the advent of shale has led to a price reduction of approximately \$20 per barrel, consistent with the long term expectations from WTI Oil futures of around $\$60-\$70$ per barrel (depending on the magnitude of the risk premia), this translates into \$131.4 billion per year in savings for oil consumers (including both household and corporate sectors). Projecting these cost savings in perpetuity (admittedly a strong assumption) and discounting them at a rather conservative rate of 10% per annum yields

[3] Another possiblity is that increased shale production might decrease uncertainty about long-run oil prices. We find some evidence for high shale returns leading to lower oil price volatility. See Appendix A.

approximately \$1.31 trillion in savings (lowering the discount rate to 5% increases this number to \$2.62 trillion). While this simple calculation is subject to many caveats, it suggests that both the impact of the shale oil technology through the supply side of the economy, as identified in our prior empirical tests, and the impact of changes in oil prices on the demand side are economically meaningful, and are of similar magnitude.

This paper proceeds as follows. First we develop a simple reduced-form asset pricing model with an explicit role for oil demand and production in Section 2. We then describe the data construction and our empirical approach in Section 3. Section 4 presents the results of our empirical analysis based on the time-series of the Shale Index. Section 5 presents the results based on the announcement events and the cross-section of industries. Section 6 concludes.

2 Model

In this section we develop a simple toy model of oil production and demand that motivates the use of asset prices to extract technology shocks.

2.1 Demand for Oil

A representative firm produces consumption goods via a Cobb-Douglas production technology

$$Y_{t+1} = A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha},$$

where A_{t+1} is an aggregate productivity shock, O_{t+1} is oil, which plays the role of an intermediate good, and K_t is capital, where the time subscript refers to the fact that capital is chosen one period ahead (i.e. before the productivity shock is realized). Capital depreciates fully after the period's production is complete. The firm acts competitively, therefore maximizing profits implies that oil prices must satisfy

$$P_t^O = (1-\alpha) A_t O_t^{-\alpha} K_t^{\alpha}$$

given the aggregate supply of oil O_t (we assume this production technology is the only source of domestic demand for oil).

2.2 Oil Supply

Total oil supply is a sum of supply generated by two oil (sub)sectors:

$$O_t = S_t^{Shale} + S_t^{Other}$$

The two sectors are:

1. shale oil, S_t^{Shale}
2. All other oil production (OPEC, Large Integrated Oil Producers, International Oil Production, etc., net of foreign demand) , S_t^{Other}

There is a continuum of competitive price-taking firms in each sector, each sharing a common, sector-specific productivity shock Z_t^i and using competitively supplied factor input L_i ('leases') at a price w_i.

Oil Company Production is given by

$$S_t^i = Z_t^i L_i^{\nu}, 0 < \nu < 1$$

Oil Company Profits

$$\Pi_t^i = P_t^O S_t^i - w_i L_i, \text{ which implies}$$

$$\Pi_t^i = P_t^O S_t^i (1 - \nu)$$

Assuming marginal cost of deploying one lease w_i is fixed, we have $\nu P_t^O Z_t^i L_i^{\nu-1} = w_i$ so that sector output is equal

$$S_t^i = Z_t^i L_i^{\nu} = \left(Z_t^i\right)^{\frac{1}{1-\nu}} \left(\frac{w_i}{\nu P_t^O}\right)^{\frac{\nu}{\nu-1}}$$

and

$$\Pi_t^i = \left(P_t^O Z_t^i\right)^{\frac{1}{1-\nu}} (1 - \nu) \left(\frac{w_i}{\nu}\right)^{\frac{\nu}{\nu-1}}.$$

The intuition behind this production function is that while the costs of drilling are roughly the same across locations, some of the drilled wells are much more productive than others and therefore are profitable to operate at lower levels of oil prices, while less productive leases are utilized only when prices are sufficiently high.

We assume that the sectors differ in their productivity Z_t^i as well as marginal cost of production w_i, which jointly determine the relative importance of each sector in total oil supply. While in general different oil sectors may differ in the degree of decreasing returns, this assumption simplifies exposition without driving any of the implications.

Assume for simplicity that one unit of capital must be invested at the beginning of the period to operate the technology, with full depreciation by the end of the period. Then returns on firms in sector i equal profits: $R_{t+1}^i = \Pi_{t+1}^i$.

We assume that all of the productivity shocks, A_t, S_t^{Shale}, and S^{Other}, together with innovations to an exogenously given stochastic discount factor M_t, are jointly lognormally distributed.

2.3 Asset Pricing

The value of capital invested in the aggregate production sector is just the present value of next period's profits:

$$V_t^i = \alpha E_t \left[M_{t+1} A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha} \right]$$

assuming no depreciation between periods ($V_t^i = K_t^i$) implies that the returns to an average firm are

$$R_{t+1}^a = \frac{\alpha A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha}}{V_t^i} = \frac{A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha}}{E_t \left[M_{t+1} A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha} \right]} = A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha-1}$$

or, in logs,

$$\begin{aligned} r_{t+1}^a &= \Delta a_{t+1} + o_{t+1} + p_{t+1} - g_A - (1-\alpha) E o_{t+1} + \alpha k_t + r_t - \frac{1}{2} Var \left[\log \left(M_{t+1} A_{t+1} O_{t+1}^{1-\alpha} K_t^{\alpha} \right) \right] \\ &= \left(E_{t+1} - E_t \right) a_{t+1} + (1-\alpha) \left(E_{t+1} - E_t \right) o_{t+1} + r_t - \frac{1}{2} \sigma_m^2 + rp^a + \frac{1}{2} \sigma_a^2 \\ &= \left(E_{t+1} - E_t \right) o_{t+1} + \left(E_{t+1} - E_t \right) p_{t+1} + r_t + rp^a - \frac{1}{2} \sigma_a^2, \end{aligned}$$

where the aggregate market equity risk premium

$$rp^a = -Cov\left(m_{t+1}, \Delta o_{t+1}\right) - Cov\left(m_{t+1}, \Delta p_{t+1}\right)$$

is assumed constant for simplicity, as is the corresponding return volatility

$$\sigma_a^2 = Var\left(\Delta o_{t+1} + \Delta p_{t+1}\right)$$

and the risk-free rate is $r_t = E_t m_{t+1} - \frac{1}{2}\sigma_m^2$.

Similarly, excess returns to oil producers in sector i are given by

$$r_{t+1}^i - r_t + \frac{1}{2}\sigma_a^2 = \frac{1}{1-\nu}\left(E_{t+1} - E_t\right) z_{t+1}^i + \frac{1}{1-\nu}\left(E_{t+1} - E_t\right) p_{t+1} + rp_t^i, \tag{1}$$

where the risk premium rp^i is determined by the conditional covariances of the shocks with the SDF innovations.

We approximate innovations to the log of total supply as

$$\begin{aligned}
\left(E_{t+1} - E_t\right) o_{t+1} &\approx \xi^{Shale}\left(E_{t+1} - E_t\right) s_{t+1}^{Shale} + \left(1 - \xi^{Shale}\right)\left(E_{t+1} - E_t\right) s_{t+1}^{Other} \\
&= \frac{1}{1-\nu}\xi^{Shale}\left(E_{t+1} - E_t\right) z_{t+1}^{Shale} \\
&\quad + \frac{1}{1-\nu}\left(1 - \xi^{Shale}\right)\left(E_{t+1} - E_t\right) z_{t+1}^{Other} - \frac{\nu}{1-\nu}\left(E_{t+1} - E_t\right) p_{t+1}
\end{aligned}$$

where $\xi^{Shale} = E\left[\frac{S_t^{Shale}}{O_t}\right]$, and we assume that Σ is a constant variance-covariance matrix of S_t^{Shale} and S_t^{Other} so that the convexity adjustment $\frac{1}{2}\left(\xi^{Shale}, 1 - \xi^{Shale}\right)\Sigma\left(\xi^{Shale}, 1 - \xi^{Shale}\right)'$ drops out.

Then aggregate market return innovations can be approximated as

$$\begin{aligned}
\left(E_{t+1} - E_t\right) r_{t+1}^a &\approx \frac{1}{1-\nu}\xi^{Shale}\left(E_{t+1} - E_t\right) z_{t+1}^{Shale} \\
&+ \frac{1}{1-\nu}(1 - \xi^{Shale})\left(E_{t+1} - E_t\right) z_{t+1}^{Other} + \frac{1-2\nu}{1-\nu}\left(E_{t+1} - E_t\right) p_{t+1}
\end{aligned} \tag{2}$$

Therefore, since the market return innovation in equation (2) is a linear combination of the sector-specific oil productivity shocks and innovations to oil prices, the long-run average share of shale oil in total oil supply, ξ^{Shale} can be identified by regressing the market excess return on the shale sector return, controlling for the other sector returns (given by (1)) as well as oil price innovations.

3 Data

Data for this project come from several sources. All data for oil production and forecasts are from the Energy Information Assocation (EIA). WTI futures returns are constructed using data from Bloomberg. Stock market data is from CRSP and Datastream. Reported revenue and analyst projections of revenue are from Thomson Reuters' IBES database. The analysis will rely heavily on two indices that we construct, one of companies with high involvement in shale oil production, and another of companies with high exposure to shale gas production. Here we explain the construction in detail.

3.1 Shale Index Construction

Shale Oil Index The objective of our index construction is to create an asset pricing measure of shale oil development. Therefore we begin with a list of all firms that may have direct shale oil exposure, that is, those firms that are SIC 1311 (Crude Petroleum and Natural Gas). We then manually collect data from the 10-Ks of these firms to assess whether a firm's assets are primarily located in areas of significant shale oil development. We exclude firms that have significant international or offshore assets, as well as firms with significant shale or non-shale natural gas assets and non-shale oil exposure. We then verify that the remaining firms have significant operating assets in the Eagle Ford Shale (TX), the Bakken Shale (ND), or the Permian Basin (TX), as these are the primary areas of shale oil development in the United States. In Table 1 we list the firms that met these criteria and report where the index components have assets.

Shale Gas Index The shale gas index was constructed in a similar manner to the shale oil index. The primary objective of our shale gas index is to have an asset pricing measure

Table 1: Construction of Shale Oil Index and Shale Gas Index

This table provides details on the components of the Shale Oil Index used in this study and Shale Gas Index used in this study. The firms in these indices are comprised of firms in SIC 1311 (Crude Petroleum and Natural Gas), that have significant asset focus on either Shale Oil or Shale Gas. Asset information was hand collected from company 10-Ks to make the determination whether a firm is shale oil or shale gas. Asset values are as of December 31, 2013.

Shale Oil Index

Ticker	Company Name	Primary Assets	Size (Assets in $ Millions)
EOG	EOG RESOURCES INC	Eagle Ford (Oil), Bakken (Oil)	30,574
PXD	PIONEER NATURAL RESOURCES CO	Permian (Oil), Eagle Ford (Oil)	12,293
CLR	CONTINENTAL RESOURCES INC	Bakken (Oil)	11,941
CXO	CONCHO RESOURCES INC	Permian (Oil)	9,591
WLL	WHITING PETROLEUM CORP	Bakken (Oil)	8,833
EGN	ENERGEN CORP	Permian (Oil)	6,622
HK	HALCON RESOURCES CORP	Bakken (Oil)	5,356
OAS	OASIS PETROLEUM INC	Bakken (Oil)	4,712
KOG	KODIAK OIL & GAS CORP	Bakken (Oil)	3,924
ROSE	ROSETTA RESOURCES INC	Bakken (Oil), Eagle Ford (Oil)	3,277
CRZO	CARRIZO OIL & GAS INC	Eagle Ford (Oil)	2,111
NOG	NORTHERN OIL & GAS INC	Bakken (Oil)	1,520
AREX	APPROACH RESOURCES INC	Permian (Oil)	1,145
CPE	CALLON PETROLEUM CO	Permian (Oil)	424
USEG	U S ENERGY CORP	Bakken (Oil), Eagle Ford (Oil)	127

Shale Gas Index

Ticker	Company Name	Primary Assets	Size (Assets in $ Millions)
CHK	CHESAPEAKE ENERGY CORP	Barnett Shale (Gas), Haynesville Shale (Gas)	41,782
RRC	RANGE RESOURCES CORP	Marcellus Shale (Gas)	7,299
COG	CABOT OIL & GAS CORP	Marcellus Shale (Gas)	4,981
XCO	EXCO RESOURCES INC	Haynesville Shale (Gas)	2,409
CRK	COMSTOCK RESOURCES INC	Haynesville Shale (Gas)	2,139
MHR	MAGNUM HUNTER RESOURCES CORP	Marcellus Shale (Gas), Utica Shale (Gas)	1,857
KWK	QUICKSILVER RESOURCES INC	Barnett Shale (Gas)	1,370
FST	FOREST OIL CORP	Haynesville Shale (Gas)	1,118
REXX	REX ENERGY CORP	Marcellus Shale (Gas), Utica Shale (Gas)	991
GDP	GOODRICH PETROLEUM CORP	Haynesville Shale (Gas)	974

of firms with a significant asset focus on shale gas. We start with the full set of firms that are SIC 1311 (Crude Petroleum and Natural Gas) and manually collect data on a firm's assets. We only include firms in our index that have assets in the major shale gas basins: Marcellus Shale (PA, WV), Barnett Shale (TX), Haynesville Shale (TX, LA), and Utica Shale (OH). Any firm whose asset focus could not be definitively categorized in these basins was excluded. Therefore, international firms, offshore firms, shale and non-shale oil firms, and non-shale natural gas firms are all excluded from this index. In Table 1 we list the firms that met the above criteria, we also report which shale gas basins firms have assets in.

4 Empirical evidence

4.1 The Shale Revolution: a Primer

Shale oil and natural gas reserves were long thought to be uneconomic to develop. For example, as recently as the late 1990s only 1% of U.S. natural gas production came from shale. Then in the early 2000s Mitchell Energy began experimenting with new techniques for drilling shale, and found that by combining horizontal drilling with hydraulic fracturing ("fracking"), natural gas from shale could be economically produced. The unlocking of shale has led to a dramatic increase in production of natural gas, which ultimately led to lower prices of natural gas in the U.S. and, consequently, electricity. With low natural gas prices and high oil prices in 2009, firms began to experiment with using shale technology to extract oil, as oil and gas are often trapped in similar geologic formations. Figure 1 displays the recent trends in oil production. Several firms were successful in adopting shale technology in oil basins, including the Permian, the Bakken formation, and the Eagle Ford shale. As Panel A shows, with the adoption of shale technology production of from these basins has increased significantly.

There are three features of the shale oil boom that make it especially interesting from an asset pricing perspective. The first is that the rise in production was unexpected, and can therefore be interpreted as a true "Technology Shock". Panel B of Figure 1 shows U.S. crude oil production from 2005 to 2014, along with monthly forecasts of future oil production from the EIA's monthly publication of Short Term Energy Outlook. Consistent with Panel

A, starting in 2012 U.S. Crude Production rises dramatically. This rise in production was unanticipated by forecasts, which consistently undershoot production for the first year of the Shale Boom, before adjusting towards the end of the period.

The second important feature of the boom is its magnitude. While clearly increased productivity is a benefit for shale oil producers, its importance for the rest of the economy hinges on the fact that this production increase is significant relative to total world supply. Panel C of Figure 1 illustrates that the increase in U.S. oil production driven by shale deposits amounts to roughly 5% of total world oil production. While this may not seem large, given the highly inelastic nature of oil demand it has a potential to have a large long-run impact on price levels. Typical estimates of long-run demand elasticity (see for instance Kilian and Murphy (2014)) are near -0.25, suggesting that a 5% increase in world supply may yield up to a 20% drop in price. While the price does not drop dramatically over the sample we consider, this period coincides with unrest in the Middle East and consequently volatile supply from the region. The recent increases in Libyan production combined with the greatly increased U.S. production have combined to depress global prices by roughly 20% in the three months since the end of our sample. Without U.S. oil production increases, it is very likely that the recent reductions in Middle East supply would have translated into significantly higher prices than those observed.

The final feature that makes this shock somewhat unique is that it originated in a small number of easily identifiable firms, which allows us to isolate innovations using stock return data.[4]. Panel D illustrates the cumulative returns to several stock price indices. The most important is the "Shale Oil Index" described in Section 3, which is designed to isolate firms which are primarily engaged in oil production in the shale oil plays. The returns to this index are plotted with several other energy producer stock indices. The first is the"Shale Gas Index", described in Section 3, the second is a "Non U.S. E&P Index", which consists of E&P firms outside of the United States. The third is an index of the four large integrated oil and gas producers on the S&P 500. The cumulative returns to the aggregate CRSP market index are also included for comparison. As Panel D shows, the shale oil firms exhibit no

[4]This is similar to the approach used by Ward (2014) to study the impact of IT innovation on aggregate stock prices

Figure 1: U.S. Oil Production and Stock Returns

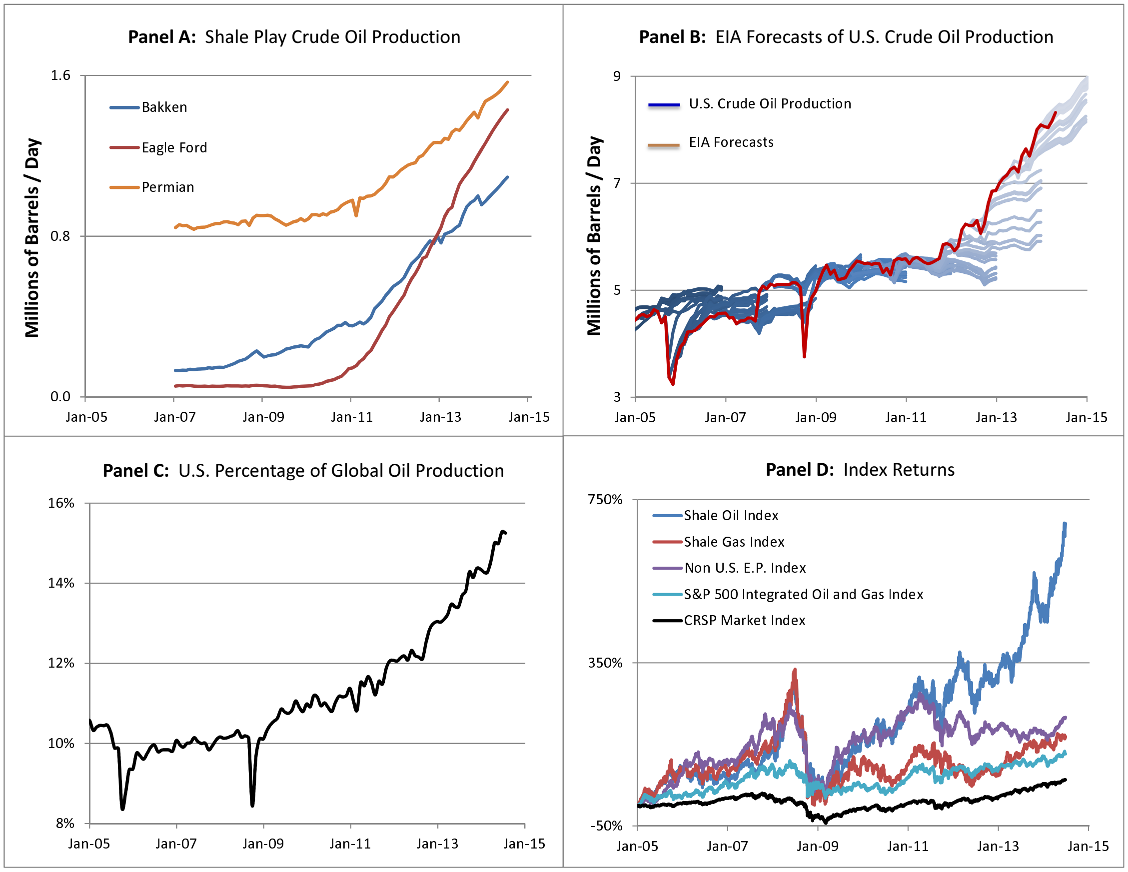

abnormal returns relative to other industry producers prior to the sharp rise in production. However, following that rise, they experience a period of extraordinary growth, rising roughly 200% in a two year time. The fact that this return was unique to these firms allows us to isolate market expectations of the value of the shock, which we do in the following section.

4.2 Constructing the Shale Shock

To isolate the impact of shale oil production we first construct daily returns on a tradable shale oil strategy by orthogonalizing the return to the shale oil index, $R^{ShaleOil}$, with respect the returns of a set of control assets. These assets are the shale gas index, $R^{ShaleGas}$; an international index of non-U.S. Exploration and Production (E&P) firms, $R^{E\&P}$; the S&P 500 index of integrated oil and natural gas producers, R^{Int}; and the return to the near month

NYMEX WTI oil future, R^{WTI}. Therefore, the resulting Shale Shock is given by

$$S_t = R_t^{ShaleOil} - \beta_{t-1}^{ShaleGas} R_t^{ShaleGas} - \beta_{t-1}^{E\&P} R_t^{E\&P} - \beta_{t-1}^{Int} R_t^{Int} - \beta_{t-1}^{WTI} R_t^{WTI} \quad (3)$$

.

The regression loadings are calculated using a rolling 1-year regression of the daily returns to $R_t^{ShaleOil}$ regressed on the set of control variables. R_t^{WTI} controls for changes in the price of oil, while $R^{E\&P}$ and R^{Int} control for general changes in the prospects for oil producers, by controling for conventional non-shale producers. To ensure that we are identifying the technological breakthroughs linked to shale oil, and not overall shale oil and natural gas, we also include $R^{ShaleGas}$ as a control. While some firms in $R^{ShaleGas}$ and R^{Int} may have some U.S. shale oil exposure, we have gone through the 10-Ks of these firms to make sure that U.S. shale oil does not represent a significant portion of existing or potential earnings for the firms in these indices. As suggested by the analysis of the toy model in 2 above, the hedged return to the shale sector can be interpreted as the pure technology shock to this sector, or more broadly can be seen as the innovation to the market participants' expectations about the ability of the shale sector to supply oil in the future.

Table 2 shows summary statistics for the constructed indices as well as the variables used in the construction. The statistics are split across two periods: Sample 1 is pre-crisis period from January 2003 to June of 2008, and Sample 2 is the post-crisis period from June of 2010 to June of 2014. We remove the crisis period so as to prevent the outlier return observations from driving any of the subsequent analysis. The start of the pre-crisis period is chosen to correspond with the beginning of the pre-crisis run-up in oil prices.

The table shows that the hedged shale strategy earns essentially no return on average over the Sample 1 (Sharpe ratio of 0.0), while subsequently earning very high returns in Sample 2 (with a Sharpe ratio of 1.03). Shale firms performed well prior to 2008, however their performance was not abnormal given the observed increase in oil prices and the high returns of comparable firms that are used as controls in the construction of the hedged strategy. In Sample 2 the situation is quite different. While the returns to all of the energy indices are once again positive, the shale strategy earns a large abnormal return over this period after controlling for the relevant industry returns. In both periods the shale strategy constructed

Table 2: Summary Statistics: Shale Shock and Index Returns

Sample 1 (01/2003 - 08/2008)

Variable	Mean	STD	Sharpe Ratio	Correlation Matrix					
				S_t	R^{Mkt}	$R^{ShaleOil}$	$R^{ShaleGas}$	$R^{E\&P}$	R^{int}
S_t	0.02	12.37	0.00						
R^{Mkt}	7.68	14.06	0.55	0.06					
$R^{ShaleOil}$	28.42	29.13	0.98	0.45	0.44				
$R^{ShaleGas}$	31.62	31.07	1.02	0.05	0.42	0.89			
$R^{E\&P}$	26.10	23.28	1.12	0.04	0.32	0.74	0.76		
R^{int}	17.15	20.86	0.82	0.04	0.63	0.76	0.75	0.66	
R^{WTI}	22.94	33.70	0.68	0.00	-0.06	0.44	0.47	0.51	0.35

Sample 2 (01/2009 - 06/2014)

Variable	Mean	STD	Sharpe Ratio	Correlation Matrix					
				S_t	R^{Mkt}	$R^{ShaleOil}$	$R^{ShaleGas}$	$R^{E\&P}$	R^{int}
S_t	15.69	15.19	1.03						
R^{Mkt}	19.44	18.71	1.04	0.15					
$R^{ShaleOil}$	19.58	36.39	0.54	0.45	0.79				
$R^{ShaleGas}$	5.28	37.73	0.14	0.03	0.74	0.85			
$R^{E\&P}$	6.43	24.03	0.27	0.06	0.74	0.79	0.77		
R^{int}	14.15	18.13	0.78	0.07	0.86	0.80	0.73	0.73	
R^{WTI}	8.24	25.71	0.32	0.06	0.48	0.63	0.55	0.60	0.56

Summary data for returns used to construct shale shock (S_t). R^{Mkt} is the cumulative return to the CRSP market index. $R^{ShaleOil}$ and $R^{ShaleGas}$ are the returns to the constructed Shale Oil and Shale Gas indices described in Table tab:index. $R^{E\&P}$ is the return to an index of Canadian Exploration and Production companies. R^{int} is the return to the S&P Integrated Oil and Gas Index. R^WTI is the return to near month NYMEX oil futures. Data are daily expressed as annualized percentage returns. S_t is constructed by orthogonalizing $R^{ShaleOil}$ with respect to the other variables (excluding the market). This orthogonalization is done using Rolling 1-year betas using daily data.

using the rolling betas has very low correlation with the control variables, suggesting that the strategy is effective in controlling for shocks impacting the energy industry as a whole.

In contrast, the correlation between the aggregate market return and the shale strategy is very different across samples, close to zero in Sample 1 (at 0.06) but noticeably higher in Sample 2, at 0.15. The average stock market returns are also much higher in the second period, at 16% per annum versus 0% over Sample 1. Discount rate shocks, such as the high but declining equity premium post-crisis, may be the main driver of this difference. However, the changing correlation between the shale strategy and the aggregate stock market suggests the possibility that positive technology shocks in the shale sector have a significant positive effect on the market as a whole. We explore this possibility in more detail below.

4.3 The Shale Shock and the Aggregate Market Returns

Having isolated the Shale Shock as a market-based measure of expectations about shale oil production, we now examine its relation with returns to the aggregate stock market. To do this we regress daily market returns on the contemporaneous value of the shock, S_t:

$$R^{Mkt} = const + \beta_S^{Mkt} S_{t+1} + \epsilon_{t+1}. \tag{4}$$

The first two columns in Panel A of Table 3 report the results over Samples 1 and 2. As the table shows, the market has a negligible loading and zero R^2 on the shock in Sample 1, but a significantly positive coefficient over Sample 2 with a slightly higher R^2. A Chow test for the difference in coefficient on the shale shock across the two samples rejects the null hypothesis that they are the same with p-value of 0.04.

The third and fourth columns of Panel A split Sample 2 into two subsamples, Sample 2a from January of 2009 to June of 2012, and Sample 2b from June of 2012 to December of 2014. Here we see that, while the coefficients are similar across the two subsamples, the last two years of the sample period exhibit much stronger explanatory power in terms of R^2. This increase in explanatory power comes from both a lower market volatility and a higher volatility of the shale shock due to a reduction in the explanatory power of the control variables for returns to the shale oil index. However, Panel B of Table 3 shows that this reduction in variance does not lead to an increase in correlation between S_t and the controls. In all periods the explanatory power of oil prices and the other energy indices for the constructed shocks is essentially zero, confirming that the hedged strategy construction is successful in eliminating these correlations.

The fact that shale oil firms become less correlated with oil prices and the return to other energy firms, while simultaneously becoming more correlated with market returns, is consistent with the emergence of a new – and volatile – shock to shale oil firms, which is also important for aggregate market returns. Additionally, the timing of this behavior is perfectly aligned with the rapid, unexpected, increase in oil output from the shale deposits exploited by these firms. These facts together suggest a causal link between the increase in shale output and aggregate market returns.

The regressions also allow us to quantify the contribution of the Shale Shock to aggregate market returns. From Table 2 we can see that market return earned an annualized 19.44% over the second sample. After controlling for the Shale Shock in the second sample, that constant in the regression reduces to 15.3%, suggesting that roughly 20% of the total market increase over this period can be attributed to shale oil growth. Given the current market capitalization of the U.S. stock market, this 22% growth translates into an approximately $2.5 trillion increase in total market value. While this is obviously an imprecise estimate, it nevertheless speaks to the large economic magnitude of these discoveries.

5 Shale News and Stock Returns

The interpretation of the prior analysis relies on the effectiveness of the constructed Shale shock as a proxy for market expectations of shale oil production. In order to address concerns that the documented correlation is the result of an omitted variable (such as discount rate variation), we now provide evidence regarding the behavior of stock returns around news announcements pertaining to shale oil and oil prices. The idea behind this identification strategy is that news announcements that are specific to shale, oil more broadly, are plausibly exogenous to other aspects of the macroeconomy, and in particular to discount rates. We implement this strategy in three slightly different but related ways. First, we instrument for

Table 3: Shale Shock and Aggregate Market Returns

	01/2003-08/2008	**01/2009-12/2014**	**01/2009-06/2012**	**06/2012-12/2014**
Panel A: Regressions of Market Return on Shale Oil Strategy				
	R^{Mkt}	R^{Mkt}	R^{Mkt}	R^{Mkt}
S_t	0.064*	0.180***	0.211***	0.145***
	(0.034)	(0.031)	(0.064)	(0.033)
Constant	0.075*	0.153**	0.115	0.205**
	(0.041)	(0.076)	(0.122)	(0.091)
Observations	1,426	1,510	860	650
R-squared	0.003	0.021	0.019	0.040
Panel B: Regressions of Shale Oil Strategy on Control Variables				
	S_t	S_t	S_t	S_t
$R^{E\&P}$	0.017	0.023	0.036	0.047
	(0.028)	(0.040)	(0.061)	(0.081)
R^{Int}	-0.007	0.036	0.055	-0.018
	(0.028)	(0.040)	(0.064)	(0.081)
$R^{ShaleGas}$	0.011	-0.037	-0.072	-0.002
	(0.023)	(0.023)	(0.054)	(0.044)
R^{WTI}	-0.009	0.020	0.029	0.026
	(0.011)	(0.023)	(0.030)	(0.044)
Constant	-0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1,426	1,510	860	650
R-squared	0.002	0.004	0.015	0.004

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Panel A shows regressions of the market return on the constructed Shale Shock (S_t) as described in Table 2 for various subperiods. Panel B shows regressions of S_t on the control variables used in the construction of the shock. Data are daily.

the time series of shale news shocks using revenue surprises around earnings announcements of major shale firms. Second, we consider the cross-section of industry returns around two major shale announcements and a significant OPEC announcement and examine the performance of this cross-section over various time periods related to shale production. Finally, in order to quantify the total economic impact of shale oil we use information in the time-series and cross-section of industry returns to estimate overall market value attributable to increases in production.

We do this by constructing a portfolio in the cross-section of industries which attempts to mimick impacts to shale production. These methods are similar to those which are standard in the asset pricing literature for quantifying risk-premium or *expected* returns, but instead we use them here to quantify *realized* returns.

5.1 Shale Oil Earnings Announcements and Aggregate Stock Returns

In order to address the issue of causality, we would like to identify exogenous shocks to shale oil firm values that can act as an instrument for returns to the Shale Index. An ideal instrument would be an announcement, or series of announcements, which provide information about shale oil production without providing material information about other important economic shocks (e.g., Savor and Wilson (2014) show that announcement dates capture the bulk of priced shocks to firm cash flows). Unfortunately, while there are announcements made by government agencies regarding oil production, they do not appear to have a material impact on the returns to oil firms, suggesting that they are not a source of new information. Instead we look at information provided by the shale oil companies' themselves as part of their regular earnings announcements, which should be private prior to the announcements as it is material to the value of the companies.

For this exercise we focus on the last two years of the sample, during which the R^2 of the market return on the Shale Shock is high and we see the largest increase in shale oil production. Though we have many companies in the Shale Index, the information released by different companies over a short time period is likely highly correlated, and therefore may become rapidly redundant. To this end, we focus on the two largest companies (in terms of shale oil assets) in the index, EOG Resources (EOG) and Pioneer Resources (PXD). To construct a measure of new information in the earnings reports, we focus on a measure of unanticipated revenue surprise, which is simply the log of the ratio of actual reported revenue to the average analyst projected revenue in the Thomson Reuters' IBES database.

We construct 15 observations, which represent announcements related to Q2 2012 to Q1 2014, with the exception of Pioneer's 2014 Q1, which is not in the IBES database. Since the earnings reports are released after market close on the announcement day, we match the revenue surprise measure to returns over the next trading day. The standard method for this analysis is a two stage least squares (2SLS) regression of R^{MKT} on $R^{ShaleOil}$, using the measure of revenue surprise as instrument for returns to the shale oil index. However, due to the well-known poor statistical properties of this procedure (especially acute in our very small sample), it may be preferable to focus on the reduced form specification of the IV regression, as suggested by Chernozhukov and Hansen (2008). Table 4 shows the results for both procedures. The OLS regressions of returns to the shale index, as well as returns to the aggregate market index, against the revenue surprise from the two firms' announcements, can be interpreted as the first stage and the reduced form specifications, respectively. Both variables show a clear positive relation with the revenue surprise of these shale firms. Even with only 15 observations, the relationship between both return variables and the revenue surprise variable is significant at the 5% level, and in fact at 1% level for the shale index return. The reduced form regression has a high R-squared of 19% for market returns on shale firms' revenue surprise. Consistent with the reduced form results, the 2SLS regression of the market excess return on the shale index return instrumented with the shale firms' surprise also recovers a strong, statistically significant relation.

As a confirmation that this relation between shale oil revenue surprise and the aggregate market return on these days is not being driven by other information revealed in the announcements, as a placebo test we repeat the analysis using the same 15 days' returns against the average revenue surprise across all firms reporting on these days. We find that there is no relation between these announcements and either shale oil returns or aggregate market returns (both the regression coefficients and the R-squared are essentially zero in

Table 4: Stock Market Returns on Shale Announcement Days

	PXD and EOG Revenue Surprises			**Market Avg. Revenue Surprises**		
Method:	OLS	OLS	2SLS	OLS	OLS	2SLS
	$R^{ShaleOil}$	R^{Mkt}	R^{Mkt}	$R^{ShaleOil}$	R^{Mkt}	R^{Mkt}
Surprise	0.213***	0.040**		0.102	-0.043	
	(0.046)	(0.017)		(0.347)	(0.123)	
$R^{ShaleOil}$			0.186**			-0.418
			(0.074)			(3.089)
Constant	0.005	0.002	0.001	0.015**	0.003	0.010
	(0.005)	(0.002)	(0.002)	(0.006)	(0.002)	(0.046)
Observations	15	15	15	15	15	15
R-squared	0.550	0.190	0.551	0.003	0.006	0.001

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table shows results of regressions of both Shale Industry and Aggregate Market returns on 15 earnings announcements for EOG Resources (EOG) and Pioneer Resources (PXD) from second quarter 2012 to third quarter 2014. For each earnings announcement a revenue surprise measure is constructed using IBES estimates and realized revenue announcements. In the first two columns this revenue surprise is then used as the independent variable in regressions of the corresponding daily return to the Shale Oil Index $R^{ShaleOil}$, and the aggregate R^{Mkt}. In the third column, the surprise is used as an instrument for $R^{ShaleOil}$ in a 2SLS regression with R^{Mkt} as the dependent variable. The last three columns repeat this analysis using the average revenue surprises from all other firms on those days as a placebo test.

all of the specifications), suggesting that information revealed in shale oil announcements is important for aggregate market returns.

5.2 Shale Exposure and the Cross-section of Realized Stock Returns

The time series of revenue surprises and market returns suggest a link between shale discoveries and the stock market. In this section we exploit heterogeneity in industry exposure to shale price innovations to quantify the impact of shale production on the market.

We use data on the 30 Fama-French industries available from Ken French's website. For the "Oil" industry, we replace the daily return with the daily return on the S&P Integrated

Oil and Gas Index in order to minimize any direct exposure to shale oil. For all industries we first construct industry market betas using data for 10 years of daily data prior to our earliest sample period. We use this out of sample construction in order to quantify changes in market exposure to shocks to shale production. For each daily industry return in the sample we then calculate the abnormal returns using this market beta. We then construct three industry specific measures of exposure to shale oil and oil prices.

The first measure we use is simply the industry's beta on the set of 15 shale revenue surprises constructed as in Section 5.1. Firms which have higher betas from the regression of daily abnormal returns on shale revenue surprises are presumably more exposed to news about shale production.

While this allows us to create a measure of exposure to shale news, revenue surprises are an imperfect proxy for news about shale oil prospects, since revenue or earnings are not necessarily the most important news about the viability of new technologies, and could also be driven by demand shocks.

Hydraulic fracturing and horizontal drilling provide the basic building blocks for shale development. However, companies need to apply this technology and then calibrate these techniques to particular oil and gas reservoirs. Often it is the case that the economics of shale in a given reservoir are unknown. Therefore when successful shale efforts are announced, significant asset revaluations occur. In many cases, a single positive well result for a reservoir can indicate the potential for hundreds of follow-on wells, which can have billions of dollars of NPV for a given company. The announcements of these positive well results represent a unique opportunity to assess how other-non-shale industries respond to unexpected announcements of significant improvements in shale supply.

The largest of these announcements in the sample is the announcement of Pioneer Natural Resources DL Hutt C #1H well in the Wolfcamp A reservoir. On July 31, 2013 after market close, Pioneer Natural Resources announced the successful test of the DL Hutt C #1H, which began production at 1,712 Barrels of Oil Equivalent per Day (BOEPD) of natural gas and crude oil, with 72% crude oil content. This was the first successful well test of the Wolfcamp A, and represented a significant improvement of shale potential across the entire Spaberry/Wolfcamp field, the worlds second largest behind only the Ghawar Field in Saudi

Arabia. Pioneers stock price increased 12.2% on this announcement, adding $2.7 Billion to the firms enterprise value.

This announcement is also the largest revenue surprise in our set, and occurs after the Shale boom was well underway. To this announcement we add the second largest revenue surprise in the set, the May 6, 2013 earnings announcement by EOG which contained substantial news about exploratory results in both the Eagleford and Bakken shale fields leading to a roughly 10% increase in EOG's stock price. We use the total abnormal industry return on these two announcement days as a proxy for industrys exposure to increases in shale productivity.

To these two proxies of shale exposure we add a measure of industry exposure to oil prices. On November 28, 2014, the OPEC released the outcome of 166th Meeting of the OPEC Conference in Vienna that occurred on the preceding day. The key result of the meeting was the decision that member countries would not cut their oil supply in response to increased supply from non-OPEC sources and falling prices. On the announcement day oil prices dropped by over 10%, and the shale index fell by roughly 8%, while the aggregate U.S. market return was essentially zero. Abnormal return on this announcement gives us a measure of exposure to an exogenous supply shock to oil prices, unrelated to technological innovation in the shale sector.

We use these three measures as explanatory values for the cross-section of industry stock returns for Samples 1, 2a, and 2b as described in Section 4.3. Figure 2 shows the results graphically while Table 5 shows the results of Fama-Macbeth regressions of industry returns on pre-sample market betas and the measures of announcement impact. Both measures of shale exposure strongly explain the cross-section of industry returns in Sample 2b, consistent with results in Section 4.3, that increased shale production had a strong impact on the stock market. The effect is also there in Sample 2a but is less pronounced and not statistically significant.

Interestingly, industry exposure to oil prices as measured by the OPEC announcement abnormal return has strong explanatory power in the cross-section in both Sample 1 and Sample 2b. Sample 1 was a period of rising oil prices and high long-run uncertainty about the oil supply. We see in this period firms which strongly benefited from the OPEC announcement

Figure 2: The Cross-section of Industry Returns and Shale News Exposures

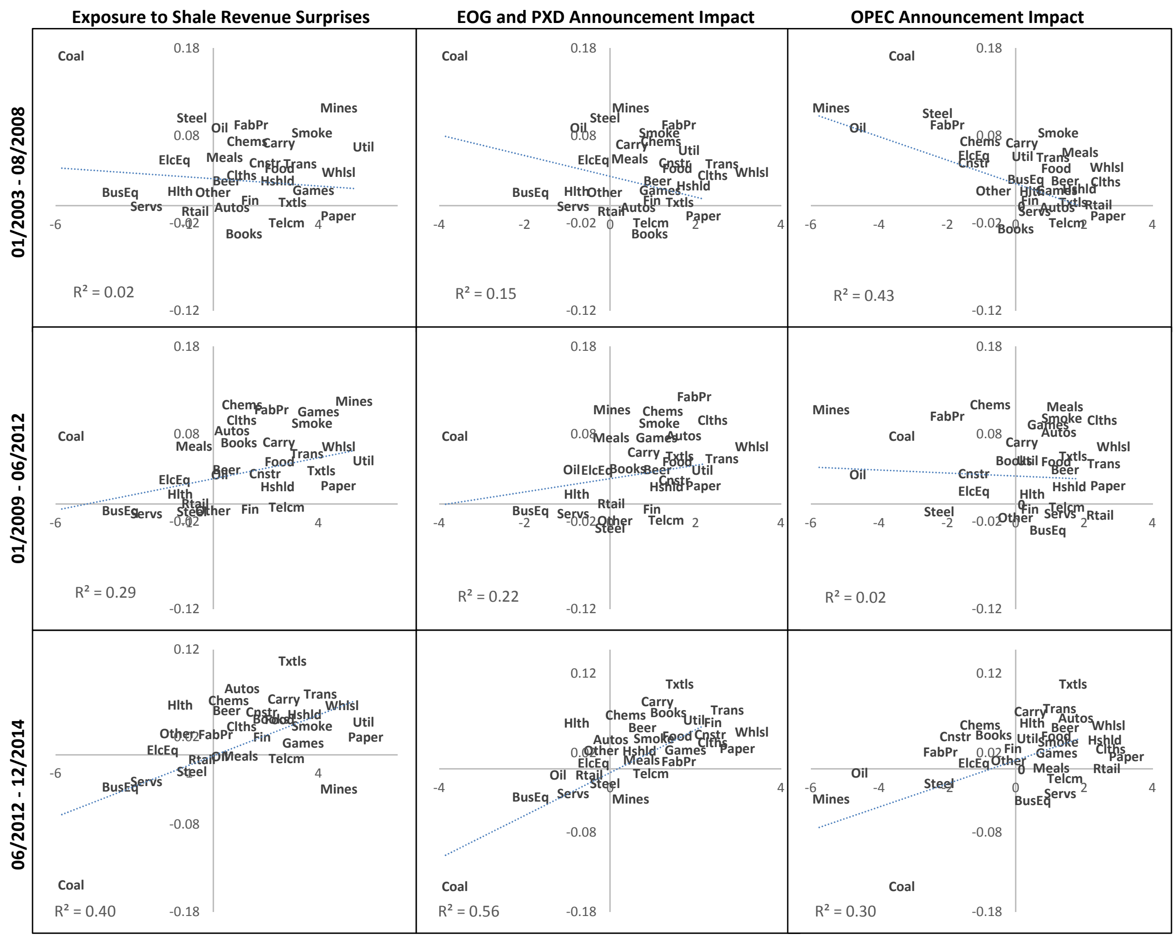

Table shows cross-section plots of industry returns on three different explanatory variables related to industry exposure to shale oil or oil prices. The first column shows the average return of each industry for each period on the Y-axis, plotted agains the Beta from a regression of industry returns on the 15 on the shale revenue surprise on the 15 shale earnings announcement days. This is analgous to the procedure used for the market return described in Table 4. The second column uses the sum of the returns on the two largest revenue surprise days in the sample. The EOG announcement of 5/7/2013 and the PXD announcemenet of 8/1/2013. The third column uses the industry return on the 11/28/2014 OPEC announcement.

have markedly lower returns. This result is novel evidence of the impact high oil prices and oil price uncertainty had on the overall equity market over this period. This effect is reversed in Sample 2b, suggesting that the rise of shale oil was a boon to industries adversely affected by high oil prices.

5.3 Constructing the Shale Factor

In this section we use the industry abnormal returns on the two major shale announcement days to construct a portfolio which has maximal exposure to shale production shocks from the cross-section of stock returns. This construction complements the Shale Shock approach in Section 4, which is motivated by the theoretical arguments in Section 2, by using the exogenous nature of the major shale announcements. Ideally we would like to construct a shale-factor-mimicking portfolio by going long industries that load positively on the shale announcement and negatively on industries that load negatively on such announcements. However, the behavior of this portfolio's returns over time would be contaminated by shocks to oil prices emanating from these industries themselves, which is presumably asymmetric in a way that drives heterogeneity in their exposure to the shale news. In order to address this concern we combine the sensitivities to the shale announcements with sensitivities to the OPEC announcement in a way that helps us isolate the direct effect of shale itself.

We construct the zero-investment factor mimicking portfolio using weights (w_i) on individual industry portfolio that are a linear functions of the total industry returns on the two large shale announcements and the return on the OPEC announcement of 11/28/2014:

$$w^i = \bar{w} + \beta_{EOG/PXD}(R^i_{5/7/2013} + R^i_{8/1/2013}) + \beta_{OPEC}(R^i_{11/28/2014})$$

The coefficients $\beta_{EOG/PXD}$ and β_{OPEC} are chosen so that the portfolio has zero return on the OPEC announcement day, and so that the largest weight in the portfolio has an absolute value of 0.4. This choice primarily effects scaling and does not significantly alter the analysis which follows. Figure 6 shows the weights for the 30 industry portfolios. We will refer to this mimicking portfolio as the "Shale Factor," and denote its returns as SF_t.

As Table 6 shows, the mimicking portfolio tends to be long firms which produce inputs

Table 5: Explaining the Cross-Section of Industry Returns with Shale News Exposures

Sample 1: 01/2003 - 09/2008

Explantory Variable	Annualized Industry Returns					
Industry Market Beta	-0.38				-1.46	-1.82
	(1.36)				(2.10)	(2.29)
Industry Response to 11/28/2014 OPEC Announcement		-5.75**			-4.16*	-5.05**
		(2.57)			(2.35)	(2.44)
Industry Response to EOG and PXD Announcements			-4.42**		-3.52	
			(1.96)		(2.28)	
Industry Beta to Shale Revenue Surprises				-2.61*		-3.20
				(1.48)		(2.25)
Constant	13.48**	13.48**	13.48**	13.48**	13.48**	13.48**
	(5.29)	(5.29)	(5.29)	(5.29)	(5.29)	(5.29)
Observations	42,780	42,780	42,780	42,780	42,780	42,780
Average R-squared	0.07	0.18	0.10	0.07	0.31	0.31
Number of Periods	1,426	1,426	1,426	1,426	1,426	1,426

Sample 2a: 01/2009 - 06/2012

Explantory Variable	Annualized Industry Returns					
Industry Market Beta	0.57				1.32	1.39
	(2.00)				(3.56)	(3.66)
Industry Response to 11/28/2014 OPEC Announcement		-0.05			-0.76	-0.40
		(3.01)			(3.30)	(3.08)
Industry Response to EOG and PXD Announcements			0.37		1.31	
			(2.75)		(3.60)	
Industry Beta to Shale Revenue Surprises				0.09		1.10
				(2.21)		(3.27)
Constant	21.20*	21.20*	21.20*	21.20*	21.20*	21.20*
	(12.86)	(12.86)	(12.86)	(12.86)	(12.86)	(12.86)
Observations	22,680	22,680	22,680	22,680	22,680	22,680
Average R-squared	0.05	0.16	0.08	0.05	0.28	0.27
Number of Periods	756	756	756	756	756	756

Sample 2b: 06/2012 - 12/2014

Explantory Variable	Annualized Industry Returns					
Industry Market Beta	2.06				4.60*	4.94**
	(1.63)				(2.51)	(2.51)
Industry Response to 11/28/2014 OPEC Announcement		8.05***			4.76*	6.57**
		(2.84)			(2.89)	(2.90)
Industry Response to EOG and PXD Announcements			6.29***		6.69***	
			(2.08)		(2.14)	
Industry Beta to Shale Revenue Surprises				3.07*		5.61***
				(1.66)		(1.99)
Constant	17.29***	17.29***	17.29***	17.29***	17.29***	17.29***
	(5.75)	(5.75)	(5.75)	(5.75)	(5.75)	(5.75)
Observations	22,620	22,620	22,620	22,620	22,620	22,620
Average R-squared	0.06	0.14	0.12	0.08	0.29	0.27
Number of Periods	754	754	754	754	754	754

Fama-Macbeth standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Cross-sectional Fama-Macbeth regressions of industry returns on market beta calculated over the pre-sample 1993 - 2003 period, shale revenue beta for each industry calculated as in Section 5.1, industry abnormal returns on two major shale announcement days, and industry abnormal returns on OPEC 11/28/2014 announcement.

Table 6: Constructing a Shale Factor

Industry	Description	Shale Factor Weight	OPEC Announcement Return (%)	EOG and PXD Announcements Return (%)
FabPr	Fabricated Products and Machinery	0.27	-2.91	3.20
Trans	Transportation	0.19	0.48	3.67
Clths	Apparel	0.16	1.32	3.70
Cnstr	Construction and Construction Materials	0.12	-2.17	2.33
Whlsl	Wholesale	0.12	-0.10	2.94
Chems	Chemicals	0.09	-1.66	2.23
Fin	Banking, Insurance, Real Estate, Trading	0.09	-0.15	2.68
Steel	Steel Works Etc	0.06	-2.89	1.68
Games	Recreation	0.05	0.34	2.57
Carry	Aircraft, ships, and railroad equipment	0.05	-0.40	2.33
Mines	Precious Metals, Non-Metallic, and Industrial Metal Mining	0.04	-5.83	0.65
Txtls	Textiles	0.04	1.28	2.77
Telcm	Communication	0.03	0.88	2.61
Books	Printing and Publishing	0.03	-0.73	2.06
ElcEq	Electrical Equipment	0.02	-1.50	1.79
Util	Utilities	-0.01	-0.04	2.03
Food	Food Products	-0.02	0.74	2.18
Autos	Automobiles and Trucks	-0.02	0.59	2.10
Hshld	Consumer Goods	-0.03	1.05	2.20
Paper	Business Supplies and Shipping Containers	-0.03	0.82	2.11
Oil	S&P Integrated Oil and Gas Producers	-0.06	-5.03	0.10
Meals	Restaraunts, Hotels, Motels	-0.06	1.14	1.94
Beer	Beer & Liquor	-0.07	0.99	1.85
Other	Everything Else	-0.07	-0.32	1.43
Smoke	Tobacco Products	-0.08	0.75	1.69
Servs	Personal and Business Services	-0.10	0.18	1.35
Rtail	Retail	-0.12	1.51	1.64
BusEq	Business Equipment	-0.12	0.07	1.19
Hlth	Healthcare, Medical Equipment, Pharmaceutical Products	-0.17	0.41	0.91
Coal	Coal	-0.40	-3.92	-2.07

First column shows weights used to construct the shale mimicking portfolio which we refer to as the "Shale Factor". Industries are the 30 Fama-French industries available from Ken French's website, with the "Oil" industry replaced by the return to the returns to the S&P Integrated Oil and Gas Producers index. The second column shows the returns to each industry on the 11/28/2014 OPEC announcement. The third column shows the sum of the returns on the two large news announcement days for EOG Resources (5/7/2013) and Pioneer Resources (8/1/2013). We construct the zero investment mimicking portfolio so that the weights (w_i) are a linear function of the total industry return on the two large shale announcements and the return on the OPEC announcement of 11/28/2014.

$$w^i = \bar{w} + \beta_{EOG/PXD}(R^i_{5/7/2013} + R^i_{8/1/2013}) + \beta_{OPEC}(R^i_{11/28/2014})$$

Where R^i_t is the return of industry i on day t. The coefficients $\beta_{EOG/PXD}$ and β_{OPEC} are chosen so that the portfolio has zero return on the OPEC announcement day, and so that the largest weight in the portfolio has an absolute value of 0.4. The procedure yields coefficients of $\bar{w} = -.28$, $\beta_{EOG/PXD} = .13$, and $\beta_{OPEC} = -.04$.

in to the production of shale oil, such as machinery, construction, and transportation firms. This suggests that the strategy is effectively targeting firms which are benefiting from the increased level of U.S. production as opposed to lower oil prices.

Table 7 shows the return to both the market and the shale mimicking portfolio over Samples 1 and 2. In Sample 1, which ends prior to the financial crisis, the shale portfolio earns a slightly negative return, while the market earns a positive return. The marginally positive exposure of the market to the shale portfolio means that the market return in the presence of zero return to the shale mimicking portfolio would have been slightly higher.

The far more striking results occur in the second period. In this period, both the market return and the shale portfolio earned highly robust positive returns. However, when the market return is regressed upon the return to the shale portfolio, the beta is high enough (roughly 0.44) to explain roughly 25% of the market return since the recovery. This is a striking result. The cross-section of industry returns constructed to coincide with the industry returns on two announcement days related to the shale industry explains a quarter of the total market increase since 2009, amounting close to $3 trillion of value. This is unlikely to be a manifestation of a cross-sectional risk premium, as the Shale Factor has negative returns in the pre-crisis sample and as the shows, and is not fully priced by the market in the second period.

As discussed earlier, this result likely understates the true impact of shale production on the U.S. economy due to the fact that decreases in oil prices or long-run oil price uncertainty, which are both consequences of the increased production, are also likely to have a positive effect but are largely excluded from our analysis. Still, it suggests that the reprieve from the high uncertainty regarding the long-run oil supply has had a significant effect on stock market returns in the last 5 years.

6 Conclusion

In a matter of a few years the technological innovations associated with fracking have revolutionized the U.S. oil market. The long run impact of this technology is uncertain, however. The continued ability of shale companies to reduce costs of extraction is actively debated,

Table 7: Explaining Market Return with a Shale Mimicking Portfolio

Sample 1: 01/2003 - 08/2008					
Average Returns		**Explanatory Regressions**			
Market Return (R_t^{Mkt})	Shale Factor (SF_t)		Market Return (R_t^{Mkt})		Shale Factor (SF_t)
7.68	-9.35*	Constant	9.79*	Constant	-10.78**
(5.99)	(5.44)		(5.89)		(5.33)
		SF_t	0.23***	R_t^{Mkt}	0.19***
			(0.03)		(0.03)
		Observations	1,426		1,426
		R-squared	0.14		0.14
Sample 2: 01/2009 - 12/2014					
Average Returns		**Explanatory Regressions**			
Market Return (R_t^{Mkt})	Shale Factor (SF_t)		Market Return (R_t^{Mkt})		Shale Factor (SF_t)
19.44***	12.32**	Constant	14.01*	Constant	8.17*
(7.49)	(5.21)		(7.15)		(4.95)
		SF_t	0.44***	R_t^{Mkt}	0.21***
			(0.05)		(0.02)
		Observations	1,491		1,491
		R-squared	0.24		0.24

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

We construct a shale mimicking portfolio in the cross-section of industry returns by going long industries with positive returns on shale announcement days and short firms with negative returns while maintainng neutrality with respect to exogenous oil price shocks using the returns to the OPEC announcement on 11/28/2014. The procedure and portfolio weights are described in Table 6. The first two columns in each panel show average annualized returns to this portfolio, which we label the "Shale Factor" (returns are denoted by SF_t), and the CRSP total market index. The last two columns show the results of two explanatory regressions, the first a regression of SF_t on the return to the aggregate market, and the second a regression of the market returns on SF_t. Data are daily from CRSP and Ken French's website.

as are the amounts of the recoverable hydrocarbons trapped in shale rock. Its importance for future economic growth also depends on the economy's long-run response to oil supply shocks, which is difficult to estimate. We use information contained in asset prices to evaluate the contribution of shale oil to the U.S. economy, to the extent that it is captured in the aggregate stock market capitalization. We find that technological shocks to shale supply capture a substantial fraction of total stock market fluctuations, suggesting that shale oil is an important contributor to the future U.S. economic growth.

References

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, August 2004, *59* (4), 1481 – 1509.

Barsky, Robert and Lutz Kilian, "Oil and the macroeconomy since the 1970s," Technical Report, National Bureau of Economic Research 2004.

Bodenstein, Martin, Luca Guerrieri, and Lutz Kilian, "Monetary policy responses to oil price fluctuations," *IMF Economic Review*, 2012, *60* (4), 470–504.

Campbell, John Y. and Tuomo Vuolteenaho, "Bad Beta, Good Beta," *American Economic Review*, 2004, *94* (5), 1249–1275.

Chernozhukov, Victor and Christian Hansen, "The reduced form: A simple approach to inference with weak instruments," *Economics Letters*, 2008, *100* (1), 68 – 71.

Christoffersen, Peter and Xuhui Nick Pan, "Oil Risk Exposure and Expected Stock Returns," *Available at SSRN 2399677*, 2014.

Cochrane, John H., "Presidential Address: Discount Rates," *The Journal of Finance*, 2011, *66* (4), 1047–1108.

Greenwald, Daniel L, Martin Lettau, and Sydney C Ludvigson, "The origins of stock market fluctuations," 2014. Working paper, National Bureau of Economic Research.

Hamilton, James D, "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, April 1983, *91* (2), 228–48.

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

Hansen, Lars Peter, John C. Heaton, and Nan Li, "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 04 2008, *116* (2), 260–302.

Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *The American Economic Review*, December 2009, *99* (3), 1053–1069.

——— **and Cheolbeom Park**, "THE IMPACT OF OIL PRICE SHOCKS ON THE U.S. STOCK MARKET*," *International Economic Review*, 2009, *50* (4), 1267–1287.

——— **and Daniel P Murphy**, "The role of inventories and speculative trading in the global market for crude oil," *Journal of Applied Econometrics*, 2014, *29* (3), 454–478.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," NBER Working Papers 17769, National Bureau of Economic Research, Inc January 2012.

Ready, Robert C, "Oil Consumption, Economic Growth, and Oil Futures: A Fundamental Alternative to Financialization," 2014. Simon Business School Working Paper No. FR 14-14.

Sadorsky, Perry, "Oil price shocks and stock market activity," *Energy Economics*, October 1999, *21* (5), 449–469.

Savor, Pavel and Mungo Wilson, "Asset pricing: A tale of two days," *Journal of Financial Economics*, 2014.

Ward, Colin, "Is the IT revolution over? An asset pricing view," 2014. University of Minnesota working paper.

A Oil Price Volatility and the Shale Shock

While the evidence presented above suggests that shale oil is important to the aggregate economy, it leaves open the question of the precise cause. Since we control explicitly for oil price changes, we think that the primary mechanism is increased economic activity through the production of Shale Oil. However, there may be additional channels through which shale oil supply affects the real economy. In this section we examine the relation between the Shale Shock and oil price volatility to see if decreasing uncertainty about oil is a potential channel for the impact of shale oil.

Recent work by Ready (2014) and Christoffersen and Pan (2014) emphasizes the importance of uncertainty in oil prices. To examine whether or not the increase in shale oil expectations has reduced uncertainty about oil, we examine the relation between the CBOE's Oil VIX measure (OVX) and the returns to our shale strategy. Rather than using daily data, we use monthly data to capture long run trends in volatility, which are potentially not captured in the day to day changes in short-term option implied volatilities used to construct the OVX.

Table 8 shows regressions of changes the OVX for 1-month periods, as well as for 3- and 12-month overlapping periods. These changes are regressed on changes in the VIX and the innovations to the Shale Strategy over the corresponding periods. Additionally, lagged values of the VIX and the OVX are included as controls. As the table shows, for both the 3-month and 12-month time horizons, innovations to the Shale Shock are strongly correlated with decreases in the OVX, controlling for changes in aggregate stock market value. This reduction in oil price uncertainty associated with rising expectations of shale producer values provides a potential channel for economic benefits from the increases in production even absent a large reduction in price.

Table 8: Shale Shock and Oil Volatility

	n = 1		n = 3		n = 12	
	$\Delta OVX_{t-n,t}$	$\Delta OVX_{t-n,t}$	$\Delta OVX_{t-n,t}$	$\Delta OVX_{t-n,t}$	$\Delta OVX_{t-n,t}$	$\Delta OVX_{t-n,t}$
$\Delta VIX_{t-n,t}$	0.482***	0.487***	0.586***	0.588***	0.709***	0.700***
	(0.062)	(0.062)	(0.069)	(0.063)	(0.096)	(0.073)
VIX_{t-n}	0.199***	0.198***	0.274**	0.267**	0.827***	0.676***
	(0.076)	(0.075)	(0.121)	(0.110)	(0.188)	(0.148)
OVX_{t-n}	-0.208***	-0.205***	-0.282**	-0.243**	-0.735***	-0.506***
	(0.074)	(0.074)	(0.119)	(0.108)	(0.218)	(0.175)
$S_{t-n,t}$		-0.273		-0.532***		-0.603***
		(0.269)		(0.204)		(0.184)
Constant	0.121	0.117	0.145	0.042	0.043	-0.243
	(0.143)	(0.141)	(0.243)	(0.222)	(0.472)	(0.366)
Observations	48	48	45	45	36	36
R-squared	0.509	0.517	0.646	0.691	0.699	0.775

Newey-West Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1