

NBER WORKING PAPER SERIES

FRACKING, DRILLING, AND ASSET PRICING:
ESTIMATING THE ECONOMIC BENEFITS OF THE SHALE REVOLUTION

Erik Gilje
Robert Ready
Nikolai Roussanov

Working Paper 22914
<http://www.nber.org/papers/w22914>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2016

We thank Ing-Haw Cheng, John Cochrane, Thomas Covert, Stefano Giglio, Wayne Ferson, Michael Johannes, Marcin Kacperczyk, Ryan Kellogg, Christopher Knittel, Leonid Kogan, Owen Lamont, Erin Mansur, Luboš Pastor, Dimitris Papanikolaou, Lasse Pedersen, Amit Seru, Andrei Shleifer, Amir Sufi, Pietro Veronesi, Ivo Welch, Mungo Wilson, Wei Xiong, Harold Zhang, and audiences at the AFA 2016 meeting, ESSFM Gerzensee, NBER (Asset Pricing and Commodity Markets meetings), Oklahoma Energy Finance Conference, Aalto, Binghamton, Cass Business School, Chicago Booth, Copenhagen Business School, Goethe University Frankfurt, Norwegian School of Management, Oxford (Said), Richmond Fed, University of Rochester, and Vienna University for valuable comments. Funding from the Jacobs Levy Equity Management Center for Quantitative Financial Research at Wharton is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Erik Gilje, Robert Ready, and Nikolai Roussanov. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Fracking, Drilling, and Asset Pricing: Estimating the Economic Benefits of the Shale Revolution
Erik Gilje, Robert Ready, and Nikolai Roussanov
NBER Working Paper No. 22914
December 2016
JEL No. G12,G13,Q43

ABSTRACT

We quantify the effect of a significant technological innovation, shale oil development, on asset prices. Using stock returns on major news announcement days allows us to link aggregate stock price fluctuations to shale technology innovations. We exploit cross-sectional variation in industry portfolio returns on days of major shale oil-related news announcements to construct a shale mimicking portfolio. This portfolio can explain a significant amount of variation in aggregate stock market returns, but only during the time period of shale oil development, which begins in 2012. Our estimates imply that \$3.5 trillion of the increase in aggregate U.S. equity market capitalization since 2012 can be explained by this mimicking portfolio. Similar portfolios based on major monetary policy announcements do not explain the positive market returns over this period. We also show that exposure to shale oil technology has significant explanatory power for the cross-section of employment growth rates of U.S. industries over this period.

Erik Gilje
The Wharton School
University of Pennsylvania
3620 Locust Walk, Suite 2400
Philadelphia, PA 19104
gilje@wharton.upenn.edu

Robert Ready
University of Rochester
533 University Ave #3
Rochester, NY, 14607
robert.ready@simon.rochester.edu

Nikolai Roussanov
University of Pennsylvania
The Wharton School, Finance Department
2400 Steinberg-Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104-6367
and NBER
nroussan@wharton.upenn.edu

1 Introduction

Technological innovations play a central role in many theoretical models of asset pricing. However, standard empirical measures of technology shocks (e.g., Solow residuals) do not appear to be large enough to explain observed movements in asset prices.¹ We provide a new approach to empirically measuring innovations in technology. We apply this approach to a sequence of shocks occurring in a particular industry with profound aggregate implications: oil. Technological innovations in shale oil development from 2012 to 2014 led to a near doubling of oil production in the U.S. and a dramatic decline in global oil prices. We find that our measure of these shale oil technology shocks explains a significant component of cross-sectional and time series variation in both asset prices and employment growth during this time period.²

Measuring the effect of a given technological innovation is empirically challenging. Typically, such innovations are difficult to observe, making it hard to trace out their impact on stock prices or real economic outcomes. A particular technological development can have diverging (often opposite) effects on different sectors of the economy. Our empirical framework uses the entire cross section of stock returns to extract innovations to latent state variables not directly observable by the econometrician. We rely on the idea that sometimes the *arrival* of relevant public news announcements is observable.³ Using the stock market reaction to the news allows us to estimate the exposures of various assets to the underlying

¹Much of the debate in empirical asset pricing 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. See, e.g. Bansal and Yaron (2004), Campbell and Vuolteenaho (2004), Hansen, Heaton and Li (2008), Cochrane (2011), Bansal, Kiku, Shaliastovich and Yaron (2014), Greenwood and Shleifer (2014), Albuquerque, Eichenbaum and Rebelo (2015), Baker, Bloom and Davis (2015), Greenwald, Lettau and Ludvigson (2014), and Campbell, Giglio, Polk and Turley (2016), for a wide range of views on the relative roles of shocks to technology, preferences, expectations, uncertainty/volatility, etc.

²Our work fits into a large literature attempting to quantify the economic impact of oil shocks on the real economy as well as asset prices, e.g. Hamilton (1983), Jones and Kaul (1996), Sadorsky (1999), Hamilton (2003), Barsky and Kilian (2004), Blanchard and Gali (2007), Dvir and Rogoff (2009), Kilian (2009), Kilian and Park (2009), Hamilton (2009), Bodenstein, Guerrieri and Kilian (2012), Ready (2016), and numerous others. Hausman and Kellogg (2015) estimate the benefits of the shale *gas* revolution on consumers. Arezki, Ramey and Sheng (2015) use large oil discoveries across countries to analyze the effect of news about *future* productivity on economic activity. Since asset prices are forward looking, the stock market should capture news about both short- and long-run productivity innovations (e.g., Beaudry and Portier (2006)).

³Our approach to empirically identifying the economic effect of technological innovations is closely related - and complementary - to recent work by Kogan, Papanikolaou, Seru and Stoffman (2012) linking news on patented technologies to equity returns.

unobservable shocks. We use key events related to shale oil development to undertake three empirical exercises designed to measure the effect of shale oil technological innovations on the economy.⁴

We first test whether an industry's exposure to shale oil technology development is linked with its stock price performance over the time period when the sequence of shale technology shocks unfolds. To identify an industry's exposure to shale oil we measure how stock prices change in response to the disclosure of a major breakthrough in shale oil extraction in the summer of 2013. This event represents the largest shale oil discovery to date, amounting to a 35% increase in expected recoverable oil reserves from the second largest oil field in the world. We trace out how different industries are affected by examining the cross-section of industry returns on this day. We find that there is significant dispersion, and that an industry's announcement return is linked with its stock performance during the shale oil time period. Specifically, a one standard deviation increase in an industry's shale discovery announcement return leads to a 4.0% higher average annual return. This relationship only exists during the time period of shale oil development (from January 2012 to September 2015), and not during earlier time periods.

In our second empirical exercise we estimate the total contribution of shale oil technology to the aggregate U.S. stock market over time, by constructing a shale mimicking portfolio based on the exposures of different industries to the shale discovery shock. The shale discovery announcement exposures can then be used to construct a factor-mimicking portfolio that tracks the unobservable innovations in shale technology over time.⁵ We use this shale-mimicking portfolio to identify the component of aggregate market fluctuations that can be attributed to shale technology shocks. Firms with high announcement returns receive a greater weight in this portfolio; firms with lower returns receive less weight. The intuition

⁴Our approach is related to several strands of asset pricing literature that focus on dates with significant public announcements. Lamont and Frazzini (2007) and Savor and Wilson (2015) focus on corporate earnings; others, such as Jones, Lamont and Lumsdaine (1998), Gürkaynak, Sack and Swanson (2005), and Savor and Wilson (2013) focus on releases of macroeconomic news; a large literature studies Federal Reserve monetary policy announcement days, e.g. Bernanke and Kuttner (2005), Savor and Wilson (2014), and Lucca and Moench (2015).

⁵The intuition for relying on cross-sectional variation to construct factor-mimicking portfolios goes back at least to Fama (1976). The approach of using asset price fluctuations to track the empirical dynamics of a hard-to-measure underlying economic variable is closely related to the economic tracking portfolios of Lamont (2001).

behind this empirical design is that there is no single asset we can use to cleanly measure innovations in shale development. However, the mimicking portfolio weights that are constructed using the slopes of the cross-sectional regressions allow us to synthetically create such an asset, building on the classic approach of Fama and MacBeth (1973). These weights are based on responses of industries' stock returns to an exogenous unexpected positive innovation in shale oil technology. We use this portfolio as an asset-price proxy for the value of shale oil development, and assess the explanatory power of this portfolio for market returns over different time periods.

We find that exposure to the shale mimicking portfolio has strong explanatory power for aggregate stock market returns from 2012 to Q3 2015 period. In total, based on the point estimate of our regressions, shale oil development is responsible for \$3.5 trillion of the increase in stock market value during this time period. The standard errors of our estimates imply that shale oil development is responsible for at least \$1 trillion of the increase in the stock market with 90% confidence.⁶ Furthermore, our shale exposure proxy has no explanatory power in earlier time periods when shale oil production was virtually nonexistent.

Finally, we assess whether the economic impact of shale oil that we measure using asset prices translates into meaningful effects on the real economy. To do this, we estimate whether the cross-section of shale discovery announcement day returns contains information about changes in industry employment. We show that the shale discovery announcement returns have significant explanatory power for the cross-section of employment growth rates of U.S. industries, indicating that the effect we identify operates through real economic channels. In the aggregate, we estimate that during the shale oil period 4,600,000 (net) new jobs are linked with the development of shale oil technology. This represents a 4.2% increase in the number of jobs across the industries in our study, compared to the aggregate number of jobs at the beginning of the shale oil period.⁷

What are the channels through which shale oil technology could affect the U.S. economy?

⁶We also estimate these magnitudes while allowing the market exposure to the shale oil mimicking portfolio to change over time by estimating rolling betas during the shale oil time period. Using this approach we obtain similar aggregate magnitudes to those in our main results.

⁷These magnitudes are consistent with estimates in the existing literature for the effects of shale gas (primarily) on employment (Hausman and Kellogg (2015), Feyrer, Mansur and Sacerdote (2015), and Allcott and Keniston (2014)). We discuss this in detail in section 3.7.

Industries' sensitivity to shale news can arise through several types of spillovers. To the extent that an increase in fracking/drilling activity increases demand for output of (imperfectly competitive) industries that provide labor or materials for shale oil extraction, the positive news about shale sector productivity is good news for these industries - we refer to this as the "supply-chain effect."⁸ To the extent that increasing income of households involved in shale oil production, directly or indirectly, improves the health of the local economies, it might benefit consumer-oriented industries that experience increasing demand for their goods - we can refer to this as the "income effect."⁹ Finally, to the extent that good news about shale oil supply can depress oil prices, it may benefit a variety of industries whose output consists of goods that are complements with oil (e.g. cars) or whose expenditure shares increase through the effect on the consumers' budget constraints - this can be called the "price effect." Additionally, a positive shock to shale extraction technology that lowers the price of oil (as well as natural gas) can have an adverse effect on industries that supply substitute energy sources, such as coal.¹⁰

We find empirical support for each of the channels described above, based on the different weights industries receive in the mimicking portfolio that we construct. For example, Oil and Gas Drilling, Business Services, Engineering Services, and Railroads receive some of the highest weights in the mimicking portfolio. All of these industries are important components in the supply chain of shale oil development. Clothes also receives a high weight, which is consistent with both an "income effect" and a "price effect." We observe that Coal Mining has the greatest negative weight in the portfolio, which is consistent with a product market rival effect driven by being a potential substitute energy source. The advantage of our methodology

⁸To the extent that shale oil development puts upward pressure on the prices of relatively less traded factors, such as labor, there might be a countervailing negative spillover effect on local firms akin to the "Dutch disease." Using detailed data on manufacturing establishments in the U.S., Allcott and Keniston (2014) also find that this effect is relatively small and that the positive supply chain and income effects dominate.

⁹Acemoglu, Finkelstein and Notowidigdo (2013) use oil reserves in the pre-shale oil period to capture shocks to local incomes. Gilje (2011) documents the impact of windfall shale oil revenues on local economies, while Cascio and Narayan (2015) focus on the increase in wages of low skilled workers and its consequences for educational attainment.

¹⁰The "price effect" is quite distinct from the others in that its magnitude can be affected by non-shale oil supply shocks. We work to separate the price effects linked to shale development, relative to price effects linked to non-U.S. oil production. To do this, we include a control for non-U.S. oil supply shocks in our regressions using a second mimicking portfolio based on an OPEC-driven oil supply shock.

that uses the entire cross section of (publicly traded) firms is that we are able to estimate the net effect of the several, often countervailing, spillover effects of a technological innovation - a major challenge in the literature (e.g., see Bloom, Schankerman and Van Reenen (2013)).

Importantly, the different channels discussed above suggest that the drop in oil prices since mid 2014 does not necessarily result in a meaningful change in the overall economic magnitudes that we estimate. Instead, it likely means that the relative importance of different channels may change. For example, while “supply chain” and “income” effects may be reduced, the “price” effect may increase. Consistent with the effect of these channels offsetting each other, while oil prices dropped 57% from their high in mid-2014 to the end of the third quarter of 2015, the level of the shale mimicking portfolio we construct remained high (despite the decline in the equity market capitalizations of the shale oil firms).¹¹

A potential concern with our methodology is that while the shale discovery announcement we rely on in our empirical design can be considered exogenous, there may have been other reasons why stock prices changed on our event day. For example, if the overall market increased for other, non-shale-oil-related, reasons, we may just be picking up the effect of high market beta as opposed to high shale exposure. We control directly for the effect of aggregate stock market beta in our main regression. However, since betas are measured with error, such controls might be imperfect. Therefore, we conduct several additional tests that demonstrate the robustness of our approach. First, we ask how likely it is that a randomly picked day would yield a cross-section of industry returns that can be successfully used to construct a factor that explains the time-series variation in both the aggregate market return (especially over the shale oil period) and returns on shale oil firms, as well as predict shale-related real activity. We find that the shale discovery announcement day is by far the single best day in the sample that can explain these features of the data. Second, instead of using the shale discovery announcement day to construct the shale mimicking portfolio, we focus on industries that are known to be part of the shale oil production chain. While this rules out some of the channels through which shale might impact various sectors of the economy, this narrower exercise still allows us to recover the underlying shale shocks and provide similar

¹¹Consistent with the economic effect we identify, shale oil production had dropped by only 8.2% from its peak.

estimates of the magnitude of their contribution to the total stock market value. Taken together, this evidence sets a high bar for alternate explanations of our results.

Nevertheless, we consider another candidate driver of aggregate stock market returns over this period, namely monetary policy. We adopt our methodology of constructing a portfolio that tracks underlying unobserved shocks by analyzing the cross-section of stock returns on the days of key announcements by the U.S. Federal Reserve (e.g., as in Savor and Wilson (2014)). We examine several sets of monetary policy announcements, including scheduled FOMC meetings or specific announcements of unconventional monetary policy. We show that such portfolios track very closely the returns on a portfolio constructed using market betas. However, while monetary policy helps explain the stock market run-up immediately following the global financial crisis in 2009, such monetary policy mimicking portfolios do not help explain any of the high market returns over the recent time period, and thus do not take any explanatory power away from the shale mimicking portfolio. This exercise, while potentially interesting in its own right, serves to highlight the general applicability of our empirical methodology, as well as the robustness of our conclusions.

This paper proceeds as follows. We describe the data, the general economic setting, and our empirical approach in Section 2. Section 3 details our econometric approach and presents the results of our empirical analysis. Section 4 presents the set of robustness tests. Section 5 concludes.

2 The Setting

2.1 The Shale Revolution: a Primer

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. Over the same time period U.S. oil production increased dramatically, from 5.4 Mb/d (million barrels of oil per day) at year end 2009 to 9.4 Mb/d at year end 2014. This increase accounted for 52.2% of overall global oil production growth. Almost all of this increase can be attributed to a breakthrough technological innovation that allows oil to be extracted from shale rock formations that were previously too costly to access. This innovation, which in-

volves a combination of two previously known technologies, 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).

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, 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, and the Eagle Ford shale. As Panel A shows, with the adoption of shale technology, production in 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 increased productivity is clearly 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, and a roughly 50% increase in production since 2009. While this may not seem large, given the highly inelastic nature of oil demand it has

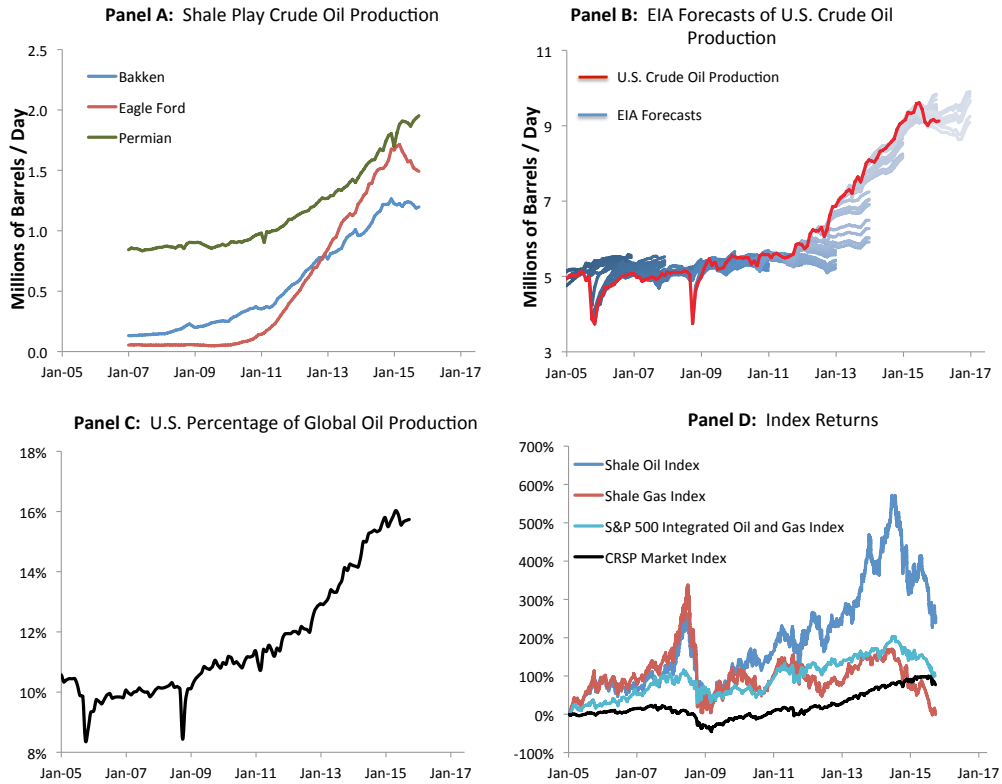
a potential to have a large long-run impact on price levels. Typical estimates of the long-run demand elasticity (see for instance Kilian and Murphy (2014)) are near -0.25, suggesting that a 5% increase in world supply may yield up to a 20% drop in price. Oil prices begin to drop in the latter part of our sample period, and then fall precipitously in the wake of the OPEC decision in November 2014, which we discuss in detail below. At the end of our sample period the level of prices is roughly half of its most recent peak.

The final feature that makes this shock somewhat unique is that it originated in a small number of easily identifiable firms which we designate as the “Shale Oil Index.” These are firms with a significant amount of production derived from shale oil. Panel D illustrates the cumulative returns of this “Shale Oil Index” to several stock price indices. The returns to the Shale Oil Index are plotted with several other energy producer stock indices. The first is the “Shale Gas Index”, described in Section 2.2, the second is a “Non U.S. E&P Index”, which consists of oil exploration and production 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 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. These stock returns are useful for understanding when asset prices began reflecting shale oil expectations. However, using a “Shale Oil Index” to precisely measure aggregate stock market effects is problematic, as discount rate shocks, as well as aggregate productivity, demand, and other shocks likely affect both the Shale Oil Index and aggregate stock prices. For this reason, we focus our main identification strategy on using asset price changes around important news announcements relevant to shale.

2.2 Data

Data for this project come from several sources. All data for oil production and forecasts are from the Energy Information Association (EIA). WTI futures returns are constructed using data from Bloomberg. Stock market data is from CRSP and Datastream (details of industry portfolio construction are in the appendix). We use NAICS code descriptions to construct 76

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



This figure plots data related to shale oil production and stock returns in different panels. Panel A plots the production of oil from the three major shale oil fields in the United States, the Bakken, the Eagle Ford, and the Permian fields. Panel B plots aggregate U.S. oil production relative to forecasts of production from the U.S. government Energy Information Administration (EIA). Panel C plots the percentage of global oil production being produced from the United States. Panel D plots the stock returns of different market composite indices, including the Shale Oil Index (defined in Appendix Table A-1), the Shale Gas Index (defined in Appendix Table A-1), the S&P Integrated Oil and Gas Index, and the CRSP Market Index. All oil production and forecast data is from the EIA and all stock price data is from the Center for Research in Security Prices (CRSP)

industry portfolios of all CRSP stocks.¹² We treat stocks of oil and gas producing companies differently, using the S&P Integrated Oil and Gas Index as our non-shale oil industry portfolio, the Shale Oil Index and the Shale Gas Index described in Appendix 3, while all the other oil producers not included in these indices populate the “Other Oil” portfolio, which is included in our main set of 76 industry portfolios.

2.3 Identification Approach: Shale News and Stock Returns

A simple model of oil production and demand presented in Appendix 1 shows that asset prices contain information about the technological shocks affecting oil production (as well as demand). It is challenging to identify these innovations empirically, since both shale oil productivity and all other shocks simultaneously drive returns to both shale oil firms and other firms in the economy (in a more general model, this would include changing discount rates, e.g. through time varying aggregate uncertainty, expectations, or preference shocks).

Our approach to overcoming this challenge involves using stock returns around news announcements pertaining to oil supply, both from shale and from non-shale sources. The idea behind this identification strategy is that news announcements that are specific to shale, and oil more broadly, are plausibly exogenous to other aspects of the macroeconomy.

We exploit heterogeneity in industry exposures to shale innovations to quantify the impact of shale production on the stock market. We consider the cross-section of industry returns around a major shale announcement and a significant OPEC Announcement and examine the performance of this cross-section over various time periods related to shale production.

2.4 Shale Discovery and OPEC Announcements

Hydraulic fracturing and horizontal drilling provide the basic building blocks for shale development. However, companies that apply these techniques typically need to calibrate them to particular oil and gas reservoirs (e.g., see Covert (2014)). Often it is the case that the economics of shale in a given reservoir are unknown ex ante. Therefore when successful shale

¹²Alternatively, one could use the standard Fama-French industries available from Ken French’s website. We construct our own industries in order to generate greater variation in exposure to oil. For example, “Airlines” are a subset of transportation industries in the most detailed Fama-French 49 industries.

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 largest of these announcements in the sample is the announcement of Pioneer Natural Resources DL Hutt C #2H 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 #2H, 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 Spraberry/Wolfcamp field, the world's second largest behind only the Ghawar Field in Saudi Arabia. Pioneer's stock price increased 12.2% on this announcement, adding \$2.7 Billion to the firm's enterprise value. The announcement of these positive well results represent a unique opportunity to assess how other firms, including in non-shale industries, respond to unexpected announcements of significant improvements in shale supply. We use industry portfolio return on this single announcement day as a proxy for an industry's exposure to increases in shale productivity.

One concern regarding our reliance on this announcement is that we might overstate the contribution of shale oil shocks to the performance of industries that are sensitive to oil prices during our sample period, since its magnitude can be affected by non-shale oil supply shocks. It is therefore important to ensure that our measure does not pick up industries' sensitivities to such price effects that are coming from other sources of oil supply. In fact, the data provides an attractive event for identifying the impact of non-shale supply shocks on oil prices. On November 28, 2014, the OPEC released the outcome of the 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. The return on this announcement day gives us a measure of exposure to an exogenous supply shock to oil prices, unrelated to technological innovation in the shale sector. Indeed, just like for the shale announcement, these returns

vary dramatically across industries.

3 Empirical evidence

3.1 Evidence from the Cross-section of Realized Stock Returns

In order to estimate the impact of shale (and oil) news on the cross section of industries we run standard Fama-MacBeth regressions of weekly excess returns of the industry portfolios on characteristics, where the latter include the shale announcement return and the OPEC Announcement return of each industry. The announcement returns are standardized to have the standard deviation equal to one. We also control for the lagged market betas of each of the industries estimated before and during the financial crisis, when we would expect shale to have a minimal impact on market returns. We use betas estimated during both these period to control for any potential change in the source of market variation during the financial crisis. We do not control for contemporaneous betas since those may be endogenous to the shale shock as industries' relative importance in the market portfolio changes:

$$r_{t+1}^j = \lambda_t^0 + \lambda_t^1 r_{ShaleDisc}^j + \lambda_t^2 r_{OPECAnn}^j + \lambda_t^3 \beta_{PreCrisis}^j + \lambda_t^4 \beta_{Crisis}^j + \epsilon_{t+1}^j. \quad (1)$$

Table 1 presents the results of these regressions across four subperiods: Pre-Crisis (01/2003 - 06/2008), Crisis (07/2008 - 06/2009), Post-Crisis (07/2009 - 12/2011), and the Shale Oil Period (01/2012 - 09/2015). Panel A presents the results using the full cross-section of industries, where as in Panel B the three key industries related to oil and gas (Shale Oil, Shale Gas, S&P Integrated producers) are excluded. Thus, all of the cross-sectional slope coefficients $\lambda_t = [\lambda_t^0, \lambda_t^1, \dots]$ are averaged over subperiods in order to understand the role of oil shock sensitivities on industry returns during the period when shale oil was – and was not – a major source of innovation.

The first result is that oil shocks are an important driver of stock returns. The effect identified through the OPEC Announcement return is strongly statistically significantly negative during the pre-crisis period of rising oil prices. The average Fama-MacBeth slope coefficient of -0.146 suggests that a one standard deviation increase in an industry's sensitivity to the

OPEC shock translates into a 14.6 basis point per week (or, about 7.6 percent per year) lower return on average over this period than an average industry. During both the crisis and the post-crisis periods the coefficient is not statistically significant, as both oil prices and stock returns fall dramatically during the crisis and then recover. Finally, during the shale period the OPEC Announcement coefficient is strongly and significantly positive at 0.142 (or 0.156 if oil firms are excluded). This is a clear manifestation of the fact that the falling oil prices during this period have lifted stock prices of firms that most benefit from low oil prices - the same firms whose valuations suffered during the period of rising oil costs before the crisis.

What is the role of shale? Unlike the OPEC Announcement, the Shale Discovery Announcement sensitivity is a significant (and positive) driver of returns only during the last period, when shale production became a significant economic force. When the shale announcement return is the only characteristic its effect is statistically significant, with a coefficient of 0.052, in the full sample, and even more strongly significant, with a coefficient of 0.117, when the shale oil, shale gas, and integrated oil and gas sectors are excluded. This suggests that the decline in oil prices driven by forces outside of the U.S. (e.g., global demand or OPEC supply) depressed valuations of U.S. shale and non-shale oil firms to a substantial degree in the most recent part of the sample. Indeed, when we control for the OPEC Announcement return the shale coefficient becomes strongly significant in both samples, with the similar magnitudes (0.077 and 0.095). Controlling for the OPEC sensitivity raises the shale slope because it allows us to disentangle two opposing effects oil prices have on U.S. firms, in their relation to the shale industry. While the “supply chain,” “income,” and “price” effects may all be positive for shale, only the direct “price effect” is positive for the OPEC shock, since it lowers oil prices without helping U.S. production. In fact the effect is negative for the firms that benefit from shale for non-price reasons, since it hurts U.S. shale oil production and therefore limits the extent of positive spillovers.

Overall, the effect of a one standard deviation increase in its sensitivity to the Shale Oil Discovery Announcement increases an industry’s stock return over the shale period by about 3 to 4 percent per annum. Controlling for the pre-crisis and crisis period stock market betas does not have any effect, suggesting that the shale announcement return is not picking up industries with (persistently) high (and low) market betas. Note that average returns

over the short subsamples that drive the Fama-MacBeth coefficients we estimate need not represent *expected* returns. The effect of shale is likely driven by a series of positive surprises - technological shocks that have a first order effect on current and future cash flows of a range of industries but may or may not change their exposure to systematic risk and expected returns.

Finally, results from these regressions also show that industry market betas are not significantly related to the cross-section of realized returns over the post-crisis or shale oil periods. This is somewhat surprising, as the both of these periods saw large positive returns to the market as a whole.¹³

3.2 Constructing the Oil Factor Portfolios

The key question we want to ask is what is the contribution of the shale oil technology shock to the variation in equity market returns over the shale oil period. Consider an economy that is subject to three types of shocks: aggregate productivity (or demand) shocks a_t , shale oil shocks z_t^{Shale} , and other shocks to oil supply, z_t^{Other} . Then the (log-linearized) returns to the aggregate equity market can be written as a sum of innovations weighted by appropriate loadings:

$$r_{t+1}^{Mkt} = E_t(r_{t+1}^{Mkt}) + \beta_a^{Mkt} (E_{t+1} - E_t) a_{t+1} + \beta_{Shale}^{Mkt} (E_{t+1} - E_t) z_{t+1}^{Shale} + \beta_{Other}^{Mkt} (E_{t+1} - E_t) z_{t+1}^{Other}$$

The toy model described in Appendix 1 presents an example of such an economy and derives this representation. We are interested in estimating the exposure of the aggregate stock market to the shale shock, β_{Shale}^{Mkt} , in particular.

While the previous analysis relies primarily on the cross-sectional variation in average returns on industries across time periods, the same identification strategy can be used to extract information about the time-series behavior of returns within each of the subsamples, and therefore shed additional light on the nature of the oil shocks that we recover. This information is contained in the time-series of the cross-sectional slopes of the Fama-MacBeth regressions. It is well known (going back to Fama (1976)) that the coefficients of the individual

¹³We discuss this finding, and its relation to monetary policy shocks, in more detail in Sections 4.2 and 4.3

Table 1: Fama-Macbeth Regression of Industry Returns on Announcement Day Returns

This table shows results from Fama-Macbeth Regressions on the cross-section of 79 weekly industry returns over different subsamples. The explanatory variables are the industry return on the Shale Discovery Announcement day (8/01/2013), the OPEC Announcement day (11/28/2014), as well as market betas calculated for both the pre-crisis and crisis periods. In Panel A all 79 industries are used, while in Panel B, three energy producer industries are excluded (Shale Gas, Shale Oil, and S&P Integrated Oil & Gas Producers). Returns are weekly.

Panel A: All Industries												
	Pre-Crisis (01/2003 - 06/2008)	Crisis (07/2008 - 06/2009)	Post-Crisis (07/2009 - 12/2011)	Shale Oil Period (01/2012 - 09/2015)		Pre-Crisis (01/2003 - 06/2008)	Crisis (07/2008 - 06/2009)	Post-Crisis (07/2009 - 12/2011)	Shale Oil Period (01/2012 - 09/2015)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Shale Discovery Ret.	0.001 [0.037]	-0.023 [-1.271]	-0.055 [-0.615]	0.034 [1.160]	-0.024 [-0.233]	0.034 [1.227]	0.034 [1.227]	0.034 [1.227]	0.052** [2.100]	0.077*** [3.123]		
OPEC Announc. Ret.		-0.146*** [-2.779]	-0.161*** [-3.303]	0.169 [0.529]	0.142 [0.455]	-0.005 [-0.152]			0.142*** [2.703]	0.147*** [2.929]		
Pre-Crisis Beta		0.194 [1.238]	0.194 [1.238]		-0.110 [-0.188]	-0.027 [-0.149]				-0.031 [-0.281]		
Crisis Beta		0.018 [0.246]	0.018 [0.246]		-0.215 [-0.206]	-0.013 [-0.069]				-0.065 [-0.858]		
Constant	0.343*** [2.785]	0.288** [2.387]	0.094 [1.065]	-0.418 [-0.399]	-0.047 [-0.085]	0.400 [1.409]	0.351 [1.237]	0.400 [1.409]	0.149 [1.177]	0.284** [2.125]	0.276** [2.318]	
Observations	21,804	21,804	21,804	3,555	3,555	10,349	10,349	10,349	14,931	14,931		14,931
Number of Weeks	276	276	276	46	46	131	131	131	189	189		189

Panel B: All Industries Excluding Shale Oil, Shale Gas, and S&P Integrated Oil and Gas												
	Pre-Crisis (01/2003 - 06/2008)	Crisis (07/2008 - 06/2009)	Post-Crisis (07/2009 - 12/2011)	Shale Oil Period (01/2012 - 09/2015)		Pre-Crisis (01/2003 - 06/2008)	Crisis (07/2008 - 06/2009)	Post-Crisis (07/2009 - 12/2011)	Shale Oil Period (01/2012 - 09/2015)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Shale Discovery Ret.	-0.063** [-2.103]	-0.047* [-1.936]	0.005 [0.026]	0.039 [0.987]	-0.010 [-0.056]	0.043 [1.199]	0.039 [0.987]	0.043 [1.199]	0.117*** [3.451]	0.095*** [3.119]		
OPEC Announc. Ret.		-0.148*** [-2.873]	-0.156*** [-3.329]	0.157 [0.482]	0.136 [0.461]	-0.005 [-0.071]			0.156*** [3.005]	0.142*** [2.941]		
Pre-Crisis Beta		0.234 [1.500]	0.234 [1.500]		-0.146 [-0.237]	-0.045 [-0.251]				-0.071 [-0.657]		
Crisis Beta		-0.011 [-0.155]	-0.011 [-0.155]		-0.188 [-0.179]	-0.001 [-0.005]				-0.035 [-0.463]		
Constant	0.424*** [3.361]	0.288** [2.381]	0.118 [1.311]	-0.492 [-0.430]	-0.057 [-0.101]	0.401 [1.410]	0.345 [1.157]	0.401 [1.410]	0.069 [0.507]	0.286** [2.137]	0.259** [2.199]	
Observations	20,976	20,976	20,976	3,420	3,420	9,956	9,956	9,956	14,364	14,364		14,364
Number of groups	276	276	276	46	46	131	131	131	189	189		189

Fama-Macbeth T-statistics in Brackets
*** p<0.01, ** p<0.05, * p<0.1

cross-sectional regressions of returns on characteristics can be interpreted as portfolio returns, since these slopes are given by

$$\lambda_t = W_t' R_{t+1}^x,$$

where $R_{t+1}^x = [r_{t+1}^1, \dots, r_{t+1}^j, \dots]$ is the vector of excess returns on the test assets and the matrix of portfolio weights is given by

$$W_t = X_t (X_t' X_t)^{-1},$$

with matrix X_t containing all of the characteristics on the right-hand side of the Fama-Macbeth regression (1), including the constant. Since $W_t' X_t = I$, the first column of W_t gives weights of a unit investment portfolio and all others correspond to zero investment portfolios that have a weighted average value of one for a given characteristic and zero for all the other characteristics. Back, Kapadia and Ostdiek (2013) refer to these as “characteristic pure play portfolios” since they are maximally diversified in the sense of minimizing the sum of squared weights across test assets, while isolating the effect of a given characteristic on the cross-section of returns by controlling for other characteristics (including betas).

Here, we start by treating the returns of industry portfolios on the shale discovery announcement day (and similarly OPEC Announcement day) as the characteristic (that remains constant over time) and use this approach to construct a trading strategy that essentially goes long industries exhibiting a positive response to the announcement day and short industries with negative return responses, while exhibiting a zero return on the other announcement day and zero market beta over the prior periods. In addition to the shale and OPEC announcement returns, we can use the pre-crisis and crisis market beta estimates as characteristics as well, constructing portfolios that capture the (potential) market rewards for exposure to beta risk. Thus, we are using the time series of individual weekly slopes λ_t that produce the Fama-MacBeth coefficients reported in the Table 1 above.

3.3 Extracting Shocks: from Cross-Section to Time Series

In order to understand the intuition behind our empirical strategy, it is useful to examine it in the context of our simple model. Consider a cross-section of N industries. Assume that the return innovation to industry $j \in [1, N]$ is given by

$$(E_{t+1} - E_t) r_{t+1}^j = \beta_a^j (E_{t+1} - E_t) a_{t+1} + \beta_{Shale}^j (E_{t+1} - E_t) z_{t+1}^{Shale} + \beta_{Other}^j (E_{t+1} - E_t) z_{t+1}^{Other} + \epsilon_{t+1}^j.$$

We want to use this cross-section of industries to construct “Characteristic Portfolios” that mimic the structural shocks. To do this we will need measures related to the exposures of industries to each fundamental shock, which is not directly observed. For estimates of exposures to the two oil productivity shocks we focus on the announcement day returns. The first day is August 1, 2013, the first trading day after the Pioneer announcement on July 31, 2013, the largest shale productivity shock in our sample. We assume that the return to industry j on this day is only driven by the shale shock (with tildes indicating innovations):

$$\tilde{r}_{ShaleDisc}^j = \beta_{Shale}^j \tilde{z}_{ShaleDisc}^{Shale}.$$

This is our key identification assumption in the sense that β_{Shale}^j is the primary source of variation in industry returns on that day (i.e., the other shocks - to aggregate non-oil productivity and non-shale oil supply - are small).

The second day is the OPEC Announcement on November 28th, 2014. We view this day as clearly having a shock to z^{Other} , but we may also allow for a possibility that this announcement signaled an increased willingness of OPEC to allow very low prices for an extended period of time, which could potentially threaten the viability of shale production in the long run, i.e. a negative shock to z^{Shale} . This yields

$$\tilde{r}_{OPECAnn}^j = \beta_{Shale}^j \tilde{z}_{OPECAnn}^{Shale} + \beta_{Other}^j \tilde{z}_{OPECAnn}^{Other}.$$

Note that we assume that the idiosyncratic shocks on these days are zero. We do this because the fundamental shocks on these days are very large, minimizing the relative importance of idiosyncratic shocks. We also assume that the other aggregate shocks are absent on the OPEC Announcement day. This is consistent with the fact that the total stock market return on the OPEC Announcement day is essentially zero, despite the fact that a number of industries clearly benefit from lower oil prices. Intuitively, the impact of the OPEC decision

on the industries that benefit from shale through the supply chain and local spill-overs is negative since the sustained OPEC supply and falling prices were expected to reduce the viability of shale production. In the aggregate, this negative effect roughly offsets the positive effect on the industries that benefit through the price channel.

We then assume that the industry-specific shocks ϵ_{t+1}^j are idiosyncratic, or at least uncorrelated with the shocks to aggregate productivity and oil productivity, or, equivalently, that market beta of an industry is completely captured by the three fundamental shocks:

$$\beta_{Mkt}^j = \frac{\beta_a^j \beta_a^{Mkt} \sigma_a^2 + \beta_{Shale}^j \beta_{Shale}^{Mkt} \sigma_{Shale}^2 + \beta_{Other}^j \beta_{Other}^{Mkt} \sigma_{Other}^2}{\sigma_{Mkt}^2} \quad (2)$$

If we focus on a period prior to the shale revolution, where we would expect the shale volatility to be zero, this simplifies to

$$\beta_{Mkt,PreShale}^j = \frac{\beta_a^j \beta_a^{Mkt} \sigma_a^2 + \beta_{Other}^j \beta_{Other}^{Mkt} \sigma_{Other}^2}{\sigma_{Mkt}^2}.$$

Now consider the standard Fama-Macbeth cross-sectional regression of industry returns on our three characteristic variables, $r_{ShaleDisc}^j$, $r_{OPECAnn}^j$, and $\hat{\beta}_{Mkt,PreShale}^j$. The slope of the regression in each period is $(X'X)^{-1} X' \bar{r}_t$, where $X = [1, \bar{r}_{ShaleDisc}, \bar{r}_{OPECAnn}, \bar{\beta}_{Mkt,PreShale}]$ is an $N \times 4$ matrix. The slope coefficient for each of the three characteristic variables at time t can be equivalently considered as the return on a portfolio where the portfolio weights are the corresponding column entries of $(X'X)^{-1} X'$. These portfolios are the maximally diversified zero investment portfolios which have a loading of one on the characteristic considered and a loading of zero on all other characteristics. Let $W = [\bar{w}_1, \bar{w}_{ShaleDisc}, \bar{w}_{OPECAnn}, \bar{w}_{MarketBeta}] = (X'X)^{-1} X'$. Thus, the Shale Discovery Portfolio has a return of one on the Shale announcement day and return of zero on the OPEC Announcement day, while the reverse is true for the OPEC Announcement portfolio. Both of these portfolios are constructed to be orthogonal to the market in the pre-shale period.

Without loss of generality we can normalize the characteristics so that $\tilde{z}_{ShaleDisc}^{Shale} = \tilde{z}_{OPECAnn}^{Other} = \beta_a^{Mkt} = 1$. The returns to the three characteristic portfolios are then given

by

$$\begin{aligned}
R_{t+1}^{ShaleDisc} &= E_t(R_{t+1}^{ShaleDisc}) + \tilde{z}_t^{Shale} + \Gamma_{ShaleDisc}^{Other} \tilde{z}_{t+1}^{Other} + \Gamma_{ShaleDisc}^a \tilde{a}_{t+1} + \bar{w}'_{ShaleDisc} \bar{\epsilon}_{t+1}, \\
R_{t+1}^{OPECAnn} &= E_t(R_{t+1}^{OPECAnn}) + \tilde{z}_{t+1}^{Other} + \Gamma_{OPECAnn}^a \tilde{a}_{t+1} + \bar{w}'_{OPECAnn} \bar{\epsilon}_{t+1}, \\
R_{t+1}^{MarketBeta} &= E_t(R_{t+1}^{MarketBeta}) + \Gamma_{MarketBeta}^a \tilde{a}_{t+1} + \bar{w}'_{MarketBeta} \bar{\epsilon}_{t+1},
\end{aligned}$$

where

$$\begin{aligned}
\Gamma_{ShaleDisc}^{Other} &= -z_{OPECAnn}^{Shale} \\
\Gamma_{ShaleDisc}^a &= \frac{z_{OPECAnn}^{Shale} \beta_{Mkt}^{Other} \sigma_{Other}^2}{\sigma_a^2} \\
\Gamma_{OPECAnn}^a &= -\frac{\beta_{Mkt}^{Other} \sigma_{Other}^2}{\sigma_a^2} \\
\Gamma_{MarketBeta}^a &= 1 + \frac{(\beta_{Mkt}^{Other})^2 \sigma_{Other}^2}{\sigma_a^2}.
\end{aligned}$$

Details are provided in Appendix 2.

If we assume that the characteristic portfolios are well diversified in the cross-section ($\bar{w}\bar{\epsilon}_t = 0$), we can identify the value β_{Mkt}^{Shale} using a regression of the market return on the three characteristic portfolios.

This method essentially takes the characteristic portfolios as functions of the fundamental shocks, and asks how much of the market return can be explained by the shale announcement characteristic portfolio after controlling for the other two portfolios, and since any idiosyncratic error is likely to bias estimates downward through a standard errors-in-variables argument, we view this as the conservative approach.

The individual values of the announcement returns and market betas, as well as the resulting portfolio weights, are reported in Appendix Table A-2. We exclude the three oil and gas indices from the portfolio construction, so that we can use the returns on these indices to validate our assumption that the shocks constructed using other industries do

indeed contain information relative to shale oil. Note that since all of the characteristic pure play portfolios are zero cost, the weights add up to one even though the characteristics do not. In particular, the industries that receive a negative weight in the Shale Discovery Portfolio do not necessarily experience a negative return on the day of the Pioneer announcement, but could simply have a weaker than average positive response (since the market return on the day was positive).

The most prominent industries in terms of their announcement return responses and portfolio weights, reported in Table A-2, are quite intuitive. Industries that receive the largest positive weights in the Shale Discovery are Oil and Gas Drilling (firms in this sector act as subcontractors for both shale and non-shale oil producers) and Business Services and Engineering Services (these two sectors include a wide range of firms, many of which are heavily involved in shale exploration and production, directly or indirectly). Railroads are also naturally sensitive to shale as the boom in oil production in the areas of the U.S. that are far from the available refining capacity or pipelines saw a dramatic rise in the shipment of oil across the country via rail. The most negative weights such as for Coal and Gold Mining are also intuitive, at least for coal, which is a major substitute for oil in heating, etc. Consumer-oriented industries, such as Clothes, receive positive weights because they have large shale announcement shocks likely due to the importance of gasoline prices in consumer budgets, as corroborated by strong positive OPEC Announcement effects of such industries. For industries like Ground Transportation there is also a clear effect of the complementarity with oil. Some industries that have strong shale announcement responses receive relatively low weights in the Shale Discovery mimicking portfolio due to the effect of controls. For example, Passenger Airlines have a well-above average Shale announcement return of 1.9 percent but receive essentially a zero weight in the portfolio because their response to the OPEC Announcement is even stronger, 5.64 percent, which is natural given the key role of fuel prices for airline profits. This industry also has a historical market beta well above one, potentially further reducing its weight in the shale portfolio. Note that the OPEC announcement returns line up very closely with the OPEC portfolio weights, loading up most on industries that benefit from low oil prices, and going short industries that benefit the most from U.S. domestic oil production, such as Oil and Gas Drilling, Mining Equipment,

Oil Pipelines, and Railroads.

3.4 Exploring the Time-series

In order to verify that the Shale Discovery Portfolio return helps us to identify the shale technology shocks that we are interested in, we begin by examining the time-series behavior of this portfolio together with the other mimicking portfolios that we constructed.


We first construct an index which reflects returns attributable to shale oil innovations \hat{z}^{Shale} by examining the residual returns to the shale discovery portfolio after controlling for the OPEC announcement portfolio and the two market beta portfolios using a time-series regression over the full sample period:

$$R_{t+1}^{ShaleDisc} = \psi_0 + \psi_1 R_{t+1}^{OPECAnn} + \psi_2 R_{t+1}^{PreCrisisBeta} + \psi_3 R_{t+1}^{CrisisBeta} + \hat{z}_{t+1}^{Shale}. \quad (3)$$

We then use this regression to construct a “Shale Mimicking Portfolio” whose returns in each week are equal to $\hat{z}_{t+1}^{Shale} = R_t^{ShaleDisc} - \psi_1 R_t^{OPECAnn} - \psi_2 R_t^{PreCrisisBeta} - \psi_3 R_t^{CrisisBeta}$. In order to verify that the cumulative return path of this mimicking portfolio is broadly consistent with the timing of shale innovations, we plot the cumulative return of this index along with measures of output and productivity from the three major shale oil plays in Figure 2. As the figure shows, the positive returns that can be attributed to the Shale Discovery announcement captured in the Fama-Macbeth regressions of Table 1 coincides with the rise of shale oil production. Starting in 2011, shale oil wells began a rapid increase, corresponding with increases in the productivity of individual wells. The number of wells leveled off in late 2012, coinciding with a pause in the rise of the shale index, which then subsequently rose again as productivity and overall output continued to increase.

While the figures are illustrative, ideally we would like to evaluate the ability of our measure of shale innovations extracted using from asset prices to track shale oil production and investment activity at a relatively high frequency. One such high-frequency measure of shale activity that we focus on is the monthly series of rig counts in the three main shale plays shown in the third panel of Figure 2. Drilling rigs represent the equipment that is used to construct and drill new shale wells. The drilling rig count as a key indicator of the capital

Figure 2: Cumulative Returns on the Shale Mimicking Portfolio



This figure plots the cumulative return on a Shale Mimicking Portfolio against various measures of productivity for the combined Bakken, Eagle Ford, and Permian shale plays. The weekly return to the Shale Mimicking Portfolio is calculated as the return to the Shale Discovery Portfolio less the returns to positions in the OPEC Announcement and Market Beta Characteristic portfolios:

$$z_{t+1}^{Shale} = R_t^{ShaleDisc} - \psi_1 R_t^{OPECAnn} - \psi_2 R_t^{PreCrisisBeta} - \psi_3 R_{t+1}^{CrisisBeta}$$

The relative positions ψ_j in the other characteristic portfolios are the slope coefficients from a single regression of the weekly Shale Discovery Portfolio returns on the returns to the other characteristic portfolios for the full 2003 to 2015 sample. The four characteristic portfolio returns are the weekly slopes of the Fama-Macbeth regressions reported in Table 1. Oil production and rig count data is from the EIA.

invested by shale firms. Table 2 shows the results of regressions during our post-2012 shale oil period, in which next month's increase in rig count ΔRC_{t+1} is the dependent variable:

$$\Delta RC_{t+1} = \delta_0 + \delta_1 \Delta RC_t + \delta_2 R_t^{ShaleDisc} + \delta_3 R_t^{OPECAnn} + \delta_4 R_t^{PreCrisisBeta} + \delta_4 R_t^{CrisisBeta} + \epsilon_{t+1}^{RC}. \quad (4)$$

As shown by the first column, there is large amount of persistence the changes in rig counts at the monthly frequency, with the current month's growth in rig counts explaining roughly 70% of the variation in the next month's. Controlling for this months change, the second column shows that this months Shale Discovery Portfolio return strongly predicts next months variation in rig counts. The remaining columns show that this relation holds controlling for the returns on the other three mimicking portfolios, which don't contribute much explanatory power. These results suggest that our constructed time series proxy is in fact strongly related to increased activity in the major shale plays.

To provide further validation that our shocks are indeed capturing information related to shale oil and other oil shocks, we examine their correlation with the major oil-related variables that were explicitly excluded from their construction: the oil price and the returns to the three oil and gas equity indices. In particular, we run the following time-series regressions

$$Y_{t+1}^j = \gamma_0^j + \gamma_1^j R_{t+1}^{ShaleDisc} + \gamma_2^j R_{t+1}^{OPECAnn} + \gamma_3^j R_{t+1}^{PreCrisisBeta} + \gamma_4^j R_{t+1}^{CrisisBeta} + \gamma_4^j R_{t+1}^{Mkt} + \omega_{t+1}^j, \quad (5)$$

where $j = \{WTI, S\&PInt, ShaleGas, ShaleOil\}$. These results are reported in Table 3. Panel A shows results from regressing the weekly WTI oil price changes on the OPEC Announcements portfolio, the Shale Discovery Portfolio, the two market beta-based portfolio and the aggregate stock market return itself. The OPEC Announcement return is extremely strongly negatively correlated with oil prices, as expected, since it is capturing the returns to firms benefitting from low oil prices and hurt by high oil prices. This result is robust across all time periods, with coefficients between -2.8 and -4.5 in magnitude. This means that a one percentage point return on the OPEC portfolio corresponds to a (roughly) three to four percent fall in the oil price. The effects of the total market return variables are not consistent over time and across specifications.

Table 2: Shale Discovery Portfolio and Shale Oil Rig Counts

This table plots the results for regressions of next month's growth in rig count in the three main shale plays (Bakken, Eagle Ford, and Permian) on this month's growth in rig count and the characteristic portfolio returns. Data are monthly from 01/2012 to 09/2015. Newey-West standard errors with one lag in parentheses.

	$\Delta RigCount_{t+1}$				
	(1)	(2)	(3)	(4)	(5)
$\Delta RigCount_t$	0.851*** [7.007]	0.860*** [7.255]	0.856*** [6.650]	0.877*** [6.864]	0.788*** [10.220]
Shale Disc. Portfolio		0.088** [2.336]		0.102** [2.423]	0.114*** [2.663]
OPEC Ann. Portfolio			-0.009 [-0.342]	-0.029 [-1.214]	-0.003 [-0.138]
Pre-Crisis Beta Portfolio					0.004 [0.492]
Crisis Beta Portfolio					0.054* [1.727]
Constant	-0.443 [-1.046]	-0.838 [-1.625]	-0.373 [-0.999]	-0.668 [-1.517]	-0.951* [-1.831]
Observations	45	45	45	45	45
R-squared	0.710	0.737	0.710	0.743	0.771

Newey-West T-Statistics in Brackets

*** p<0.01, ** p<0.05, * p<0.1

The coefficient of the Shale Discovery Portfolio is positive and statistically significant only in the recent shale oil period, with a positive shale return of 1% corresponding to around a 2 percentage point rise in the oil price. This positive coefficient suggests that the Shale Discovery portfolio is primarily driven by industries that benefit from the positive spillovers generated by the shale oil production, more so than by firms benefitting from a potential effect of shale on the oil price. This validates our use of the OPEC Announcement as a control for non-U.S. oil supply that drives much of the variation in the price of oil. Indeed, the R^2 of almost all of these regressions are between 40 and 60 percent, with most of the explanatory power coming from the OPEC Announcement returns. A notable exception is the post-crisis period, when the U.S. equity market return captures a large part of variation in the WTI price, suggesting an important role for the strong demand for oil driving up the prices as the economy recovers from the Great Recession.

Panel B presents results from regressing the S&P Integrated Oil & Gas Index returns on the same variables. The evidence here is similar, as the OPEC Announcement portfolio is picking up the variation in the oil prices, which drives much of the fluctuations in the oil firm returns. The Shale Discovery portfolio is positively correlated with the integrated producers' returns during the shale period, but not after controlling for the market return, when the effect becomes insignificant - and marginally negative in the recent period). Panel C presents similar evidence for the Shale Gas index, suggesting that while shale oil and gas might benefit from the same forces that increase global oil prices, there is not a particularly strong direct connection between the two during the shale oil period that we focus on.

Finally, Panel D shows the same regressions for the Shale Oil Index. Here the effect of the Shale Discovery Portfolio is markedly different, even though the OPEC Announcement effect is very similar to those above. The two shale-related portfolio returns are extremely strongly correlated during the shale period, with coefficients between 2 and 4, approximately (the smaller coefficient when controlling for the market return). During the other time periods the correlation is much weaker and not robustly significant, as expected. This suggests that, even though the Shale Discovery Portfolio return explicitly does not include any shale oil firms, it loads strongly on industries that benefit from the shale revolution.

Table 3: Explaining Oil Prices and Index Returns with Characteristic Portfolio Returns

This table plots time series regressions where log oil price changes and returns to three energy producer indices as the dependent variables. The Shale Oil Index is constructed from oil and gas E&P firms primarily involved in shale oil extraction. The Shale Gas Index is constructed from E&P firms primarily involved in shale gas extraction. The S&P Oil and Gas Producer Index is comprised of energy majors in the S&P 500 index. Weekly log oil price changes and the index returns are regressed against returns to characteristic portfolios. The characteristic portfolios are the weekly slope coefficients from cross-sectional regressions of 76 industry returns (not including the three producer indices) on the industry Shale Discovery Day (8/1/2013) return, the OPEC Announcement day (11/28/2014) return, the industry market beta estimated in the Pre-Crisis period (01/2003 - 06/2008), and the industry market beta estimated in the Crisis period (07/2008 - 06/2009). In all specifications the original cross-sectional regressions are done using all four independent variables. The time-series regressions are conducted both with and without the Shale Discovery Portfolio return in each of our four subperiods. Data are weekly.

Panel A: Oil Price Change								
	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Portfolio	-0.967	-0.964	-1.522	-0.017	-0.616	-0.826	2.293***	1.225**
	[-1.284]	[-1.299]	[-0.651]	[-0.008]	[-0.614]	[-1.282]	[3.829]	[2.314]
OPEC Announc. Portfolio	-2.892***	-2.875***	-3.297**	-3.912***	-3.271***	-3.102***	-4.404***	-4.072***
	[-8.908]	[-8.721]	[-2.157]	[-2.799]	[-4.578]	[-5.746]	[-8.679]	[-8.516]
Pre-Crisis Beta Portfolio	-0.311***	-0.213	-0.924**	-0.975***	-0.340	-0.342**	-0.216	-0.410***
	[-3.300]	[-1.430]	[-2.218]	[-2.626]	[-1.198]	[-1.970]	[-1.434]	[-2.848]
Crisis Beta Portfolio	-0.138	-0.094	0.453***	-0.175	0.117	0.040	0.175	0.172
	[-0.801]	[-0.536]	[3.006]	[-0.552]	[0.421]	[0.228]	[0.764]	[0.796]
Market Return		-0.187		0.853**		1.054***		0.652***
		[-0.793]		[2.190]		[12.300]		[5.157]
Constant	0.107	0.114	-1.509	-0.973	0.449	0.060	-0.045	-0.183
	[0.440]	[0.480]	[-1.090]	[-0.749]	[1.189]	[0.227]	[-0.184]	[-0.767]
Observations	276	276	45	45	131	131	189	189
R-squared	0.301	0.304	0.390	0.446	0.184	0.605	0.435	0.503

Panel B: S&P Integrated Oil & Gas Index								
	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Portfolio	0.430	0.243	-2.341**	-1.794***	0.914	0.121	1.747***	-0.103
	[1.144]	[0.987]	[-2.556]	[-3.154]	[1.579]	[0.457]	[4.423]	[-0.464]
OPEC Announc. Portfolio	-1.806***	-2.029***	-0.442	-0.634*	-1.412***	-1.328***	-2.165***	-1.727***
	[-10.728]	[-18.033]	[-0.904]	[-1.806]	[-5.109]	[-7.918]	[-8.287]	[-10.906]
Pre-Crisis Beta Portfolio	0.316***	-0.334***	-0.267	-0.404***	0.410***	-0.449***	0.062	-0.442***
	[5.582]	[-6.692]	[-1.593]	[-3.956]	[3.711]	[-5.268]	[0.716]	[-7.018]
Crisis Beta Portfolio	0.740***	0.171**	0.702***	0.060	0.755***	-0.333***	0.581***	0.021
	[6.097]	[2.216]	[11.817]	[0.735]	[6.817]	[-3.712]	[5.303]	[0.276]
Market Return		1.094***		0.784***		1.267***		1.072***
		[16.807]		[7.790]		[15.303]		[18.084]
Constant	0.137	0.063	-0.451	-0.127	0.343*	-0.115	0.147	-0.095
	[1.165]	[0.746]	[-1.036]	[-0.389]	[1.906]	[-1.182]	[1.024]	[-1.110]
Observations	276	276	45	45	131	131	189	189
R-squared	0.465	0.711	0.775	0.887	0.597	0.885	0.432	0.799

Table 3: Explaining Oil Prices and Index Returns with Characteristic Portfolio Returns (Continued)

Panel C: Shale Gas Index								
	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Portfolio	0.549	0.340	2.139	3.215**	0.074	-0.579	1.508**	-0.556
	[0.963]	[0.793]	[1.488]	[2.065]	[0.085]	[-0.783]	[2.331]	[-1.008]
OPEC Announc. Portfolio	-3.538***	-3.788***	-5.520***	-5.898***	-2.977***	-2.908***	-3.663***	-3.174***
	[-13.485]	[-15.380]	[-8.180]	[-12.234]	[-5.828]	[-6.003]	[-9.077]	[-10.095]
Pre-Crisis Beta Portfolio	0.309***	-0.420***	-0.303	-0.574***	0.144	-0.563***	-0.074	-0.636***
	[3.840]	[-4.617]	[-1.110]	[-4.477]	[0.800]	[-2.699]	[-0.406]	[-3.525]
Crisis Beta Portfolio	0.936***	0.298**	1.487***	0.223	1.729***	0.832***	1.232***	0.608**
	[4.873]	[2.082]	[8.153]	[0.782]	[12.579]	[3.264]	[5.580]	[2.511]
Market Return		1.225***		1.544***		1.044***		1.196***
		[10.016]		[5.645]		[5.250]		[6.596]
Constant	0.262	0.179	-0.253	0.384	0.326	-0.052	0.152	-0.119
	[1.534]	[1.234]	[-0.315]	[0.776]	[1.107]	[-0.189]	[0.624]	[-0.549]
Observations	276	276	45	45	131	131	189	189
R-squared	0.553	0.676	0.838	0.926	0.676	0.738	0.419	0.568

Panel D: Shale Oil Index								
	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Portfolio	1.014**	0.824**	-0.305	0.649	1.373**	0.670	4.019***	1.987***
	[2.108]	[2.271]	[-0.205]	[0.859]	[1.967]	[1.226]	[7.666]	[4.366]
OPEC Announc. Portfolio	-3.464***	-3.691***	-3.298***	-3.633***	-3.170***	-3.096***	-4.179***	-3.698***
	[-14.748]	[-16.656]	[-4.924]	[-9.283]	[-6.916]	[-8.110]	[-11.716]	[-12.642]
Pre-Crisis Beta Portfolio	0.175**	-0.487***	-0.249	-0.489***	0.271*	-0.491***	0.306**	-0.248**
	[2.516]	[-6.594]	[-1.325]	[-4.642]	[1.723]	[-2.758]	[2.128]	[-2.090]
Crisis Beta Portfolio	0.732***	0.152	1.231***	0.111	1.558***	0.593***	1.245***	0.631***
	[4.226]	[1.205]	[11.971]	[0.884]	[15.121]	[3.321]	[6.524]	[3.750]
Market Return		1.113***		1.369***		1.124***		1.178***
		[10.782]		[10.520]		[6.779]		[12.830]
Constant	0.175	0.099	-0.434	0.131	0.534**	0.128	0.450**	0.184
	[1.165]	[0.782]	[-0.688]	[0.333]	[2.131]	[0.594]	[2.243]	[1.243]
Observations	276	276	45	45	131	131	189	189
R-squared	0.561	0.686	0.838	0.946	0.718	0.800	0.608	0.756

T-Statistics in Brackets
*** p<0.01, ** p<0.05, * p<0.1

3.5 Explaining Stock Market Performance

Ultimately, we would like to understand the role of the technological innovations in the shale oil sector on the U.S. stock market as a whole. A natural way to do this is via performance attribution, which, in our case, amounts to regressing the market return on the same portfolios we used to correlate with the oil price and oil and gas indices above:

$$R_{t+1}^{Mkt} = \gamma_0^{Mkt} + \gamma_1^{Mkt} R_{t+1}^{ShaleDisc} + \gamma_2^{Mkt} R_{t+1}^{OPECAnn} + \gamma_3^{Mkt} R_{t+1}^{PreCrisisBeta} + \gamma_4^{Mkt} R_{t+1}^{CrisisBeta} + \omega_{t+1}^j, \quad (6)$$

Table 4 presents the results.

In Panel A, we regress the market return on only the two announcement day characteristic portfolios. Since the pre-crisis and crisis betas are included in the Fama-Macbeth regressions, the correlation of these two portfolios to the market return is zero by construction in these two periods, as is shown in the first two columns. In the second two columns, this is no longer the case. However, in the post-crisis period we see that the Shale Discovery Portfolio still has very little explanatory power for the market, while the OPEC portfolio is now very negatively correlated with the market, due to the fact that in this period the aggregate market returns are much more positively correlated with oil prices. The more interesting results come in the shale oil period. In this period, which saw high returns to both the shale portfolio and the market, we also see a large significant exposure of the market to the shale portfolio. Including the shale portfolio in a regression leads to a 15% increase in R^2 .

Panel B repeats this analysis, but this time including the two market beta characteristic portfolios. Prior to the crisis we see insignificant positive exposure of the market to the Shale Discovery Portfolio, suggesting that it has little explanatory power for the market in these periods, although this is largely by construction. In the post-crisis and shale periods, we see that our pre-crisis and crisis beta portfolios exhibit large positive correlations with the market. In particular, these portfolios explain 70% of the variation in market returns during the post-crisis period, and essentially drive out the explanatory power of the OPEC Announcement Portfolio in this period. Again in this period, we see very little impact of the Shale Discovery Portfolio on the market.

The most striking results again occur in the shale period. In this period, while the two

market beta portfolios are still significantly correlated with the market return, they no longer explain as much of the total variation in the market. When the Shale Discovery Portfolio is included in the regression, the beta of the market on shale is again much higher (roughly 1.73) and highly statistically significant. Moreover, adding the Shale Discovery Portfolio to the regression increases the R^2 from 0.25 to 0.40, suggesting that during this period news about shale oil is responsible for about 15% of the variation in the aggregate stock market. In the other periods the contribution of shale to the market variance is essentially zero.

3.6 Economic Magnitudes

We can use the coefficients in Table 4 to estimate the overall value effect of shale oil development. The last row of each panel in Table 4 gives the change in the constant term in the regression of the market return on the characteristic portfolios that is created by including the shale portfolio, $\Delta\gamma_0^{Mkt} = \gamma_0^{Mkt} - \gamma_0^{Mkt}|_{\gamma_1^{Mkt}=0}$. Here the restricted model ($\gamma_1^{Mkt} = 0$) is estimated jointly with the unrestricted one, and standard errors for the difference are obtained via GMM. In the full regression including the beta controls (i.e., as in Panel B), this value is 11.6 basis points for the shale period. Therefore, over the 189 week shale oil period, the total cumulative return is 11.6 basis points \times 189 = 21.92%. As a robustness exercise we allow for time-varying exposures of the market return to the mimicking portfolios by estimating rolling regressions and obtain similar results (these are reported in the Appendix).

The estimates above imply that the overall value effect of shale, implied by asset prices is 21.92% of the U.S. total equity market capitalization as of the beginning of the shale period. The total market value at the beginning of the shale period was \$16 trillion, therefore the total value effect derived from our methodology is 21.92% \times \$16 trillion = \$3.5 trillion. The standard error on this point estimate does suggest a range of economic magnitudes. To be conservative, one could use the standard errors of our estimates and calculate a lower bound: for example, our estimates imply that shale oil development is responsible for at least \$1 trillion of the increase in the stock market with 90% confidence.

How plausible are these figures? As a simple check, we can compare this figure to an estimated value of the capital expenditures in shale oil extraction over time. According to the Oil & Gas Journal, capital spending by the oil and gas Industry in the U.S. was estimated

Table 4: Explaining Market Returns with Characteristic Portfolio Returns

This table shows time series regressions of aggregate stock market returns on characteristic portfolio returns in four subperiods. The characteristic portfolio returns are constructed as the weekly slope coefficients in a Fama-Macbeth regression of the cross-section of industry returns on the OPEC Announcement Return, the Shale Discovery Return, and industry market betas calculated in both the pre-crisis and crisis periods. The three oil indices are not included in the original cross-sectional regressions. Panel A shows regressions of market returns on the two announcement day characteristic portfolios. The exposure of the market to these two portfolios are zero by construction in the pre-crisis and crisis periods. Panel B repeats the exercise but this time including all four characteristic portfolios.

Panel A: No Market Beta Characteristic Portfolios								
	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Portfolio		0.000 [0.000]		0.000 [0.000]		1.107* [1.699]		1.749*** [4.768]
OPEC Announc. Portfolio	0.000 [0.000]	-0.000 [-0.000]	-0.000 [-0.000]	-0.000 [-0.000]	-1.572*** [-4.661]	-1.780*** [-5.026]	-0.521** [-2.347]	-0.862*** [-3.024]
Constant	0.161 [1.394]	0.161 [1.392]	-0.553 [-0.581]	-0.553 [-0.586]	0.345 [1.476]	0.296 [1.269]	0.355*** [2.773]	0.237* [1.883]
Observations	276	276	46	46	131	131	189	189
R-squared	0.000	0.000	0.000	0.000	0.146	0.167	0.037	0.187
Market Return Explained by Shale Portfolio								
Change in Intercept						0.049 [0.992]		0.118** [2.144]
Panel B: With Market Beta Characteristic Portfolios								
	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Portfolio		0.171 [0.602]		-0.697 [-0.975]		0.626* [1.734]		1.726*** [5.738]
OPEC Announc. Portfolio	0.248** [2.426]	0.204* [1.939]	-0.093 [-0.485]	0.245 [0.676]	0.076 [0.371]	-0.066 [-0.318]	-0.067 [-0.357]	-0.408* [-1.957]
Pre-Crisis Beta Portfolio	0.594*** [14.056]	0.595*** [13.504]	0.189 [1.442]	0.175 [1.312]	0.692*** [9.882]	0.678*** [9.594]	0.475*** [6.628]	0.470*** [6.445]
Crisis Beta Portfolio	0.514*** [5.547]	0.521*** [5.126]	0.815*** [11.830]	0.819*** [11.858]	0.861*** [11.813]	0.859*** [12.068]	0.526*** [5.076]	0.522*** [5.984]
Constant	0.067 [0.900]	0.068 [0.916]	-0.359 [-0.903]	-0.413 [-1.072]	0.390*** [3.221]	0.362*** [2.998]	0.342*** [2.985]	0.226** [2.075]
Observations	276	276	45	45	131	131	189	189
R-squared	0.594	0.595	0.815	0.819	0.769	0.776	0.249	0.395
Market Return Explained by Shale Portfolio								
Change in Intercept		-0.001 [-0.244]		0.054 [0.668]		0.029 [1.138]		0.116** [2.211]

T-Statistics in Brackets
 *** p<0.01, ** p<0.05, * p<0.1

to be \$338 billion in 2014. The Baker Hughes rig count statistics imply that roughly 78% of this activity is associated with shale oil development. Despite the recent downturn in prices, the EIA expects shale oil development to persist for many years. Assuming a 15 year life on this development and a 10% annual discount rate suggests that the present value of cash flows associated with shale oil development is \$2 trillion. However, the 15 year life assumption above is based on existing shale oil production relative to proved reserves, as outlined by the EIA, and does not include probable and possible reserves. The extent to which these additional reserves are produced or new discoveries are made, the higher the expected life of the development will be and the greater the value of the resource. This back of the envelope calculation is consistent with the \$3.5 trillion magnitude implied by asset prices.

Moreover, our method does not distinguish between the market impact of the reductions in oil prices and a decrease in long-run oil supply uncertainty, which likely resulted from the emergence of shale oil. Given the potentially quite high counterfactual levels of oil prices in the absence of shale, as well as the size of the shale industry, both of these factors would be consistent with a large effect on asset prices.

3.7 Shale Discovery Announcement Returns and Employment Growth

So far we have documented a substantial effect of shale oil on equity market values. We also assess whether the shale oil technology shock is channeled through real activity. In order to verify that this is indeed the case we examine employment growth over our sample period at the level of industries that were used in our industry portfolio construction. We build a detailed dataset of month-by-month employment by industry from the Bureau of Labor Statistics, and then calculate the aggregate growth in different industries across the time periods we focus on in our study. In Table 5 we report the results of regressions where we estimate the effect of the return from the shale discovery announcement day on average annual employment growth during different time periods ($\Delta Empl_t^j$).

$$\Delta Empl_t^j = \lambda_E^0 + \lambda_E^1 r_{ShaleDisc}^j + \lambda_E^2 r_{OPECAnn}^j + \lambda_E^3 \beta_{PreCrisis}^j + \lambda_E^4 \beta_{Crisis}^j + \epsilon_t^j. \quad (7)$$

As can be seen from the results in Table 5 there is a positive and statistically significant coefficient on the announcement return during the shale oil period. The economic interpretation of the coefficient in specification (7) of Table 5 is that if an industry's return on the shale discovery announcement day is one standard deviation higher, it experiences a 0.56% increase in average annual employment growth over the shale oil period (the announcement returns are not standardized by standard deviation, this estimate is based on the point estimate of 0.724 multiplied by the sample standard deviation of 0.77). To estimate aggregate employment effects we multiply an industry's announcement return by the point estimate of 0.724 in Table 5 and the number of years in the shale oil period, then scale an industry's 2012 employment by this estimate. We then sum up this net employment change across all industries, and find that overall employment increased by 4,600,000 jobs, due to industry exposures to shale technology. This figure is a 4.2% increase in the number of jobs, in aggregate, across the industries in our study during the shale oil time period. When we control for the OPEC Announcement return and industry betas, the coefficient increases to 0.928. However, when using this coefficient to construct aggregate job estimates to compare to actual job growth one needs to consider that the OPEC Announcement effect partially offsets the shale announcement effect for some some industries. Therefore the coefficient in specification (8) on shale announcement, is offset by job losses due to the OPEC Announcement when compared to actual changes in employment. Using the coefficient on shale discovery returns in specification (8) of Table 5, we calculate an aggregate employment effect linked to shale technology of 5,800,000 jobs. It is also important to note the standard errors on our point estimates. Using these we can construct a lower bound estimate for the effect of shale oil technology on employment using specification (8) of Table 5, and find that at least 1,600,000 new jobs are linked with shale oil development with 90% confidence.

The magnitudes we identify above are broadly consistent with recent studies that measure both direct and indirect effects of shale development on employment, mostly focusing on natural gas extraction (in contrast to our focus on shale oil). For example, Hausman and Kellogg (2015) find the indirect effects of shale gas on manufacturing employment alone of around 280,000 jobs, or a 2.6% increase (manufacturing accounts for roughly 10% of the overall private sector employment in our sample). Alternatively, Feyrer et al. (2015) estimate

that the direct effects of shale gas and oil development contributed about 640,000 jobs through 2012 in the areas of shale discoveries. Allcott and Keniston (2014) find that counties with high resource endowments experience employment growth of 2.87% when employment in resource extraction doubles overall. Using recent prices and production data from the U.S. Energy Information Administration, we estimate the dollar value of shale oil production to be 2.5 times larger than that for shale gas production. To the extent this ratio is a rough proxy for the relative economic impact of shale oil versus shale gas, combining the indirect and direct estimates above, and applying the appropriate scaling yields values well within the confidence interval of our point estimates for aggregate employment growth.

As a falsification, we show that during earlier, non-shale oil time periods, there is no statistically significant relationship between the return an industry experiences on the shale discovery announcement day and an industry's employment growth. Taken together, the evidence presented in Table 5 suggests that shale not only influenced asset prices, but had important real effects on the economy.

Table 5: Industry Shale Exposure and Employment Growth

This table reports regressions of employment growth on the shale discovery return. We aggregate up employment growth over each of the different time periods of our study: pre-crisis, crisis, post-crisis, and shale oil. Therefore, the unit of observation in these regressions is at the time period-industry level. Each time period is normalized to reflect the average annual employment growth during that time period. Data on employment was collected from the Bureau of Labor Statistics.

	Pre-Crisis		Crisis		Post-Crisis		Shale Oil Period	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Shale Discovery Return	-0.157 [-0.28]	0.196 [0.39]	-1.288 [-1.26]	-1.026 [-1.09]	-0.704 [-1.10]	-0.371 [-0.63]	0.724* [1.70]	0.928** [2.19]
OPEC Announcement Return		-0.623*** [-3.85]		0.260 [0.86]		-0.807*** [-4.24]		-0.253* [-1.85]
Pre-Crisis Beta		-0.005 [-1.10]		-0.015 [-1.66]		0.007 [1.34]		-0.005 [-1.22]
Crisis Beta		0.004 [0.77]		-0.019** [-2.07]		-0.003 [-0.57]		-0.001 [-0.15]
Constant	0.009 [0.93]	0.008 [0.47]	-0.062*** [-3.65]	0.053 [1.61]	0.011 [1.03]	-0.022 [-1.06]	0.002 [0.28]	0.019 [1.27]
R-squared	0.001	0.242	0.021	0.235	0.016	0.238	0.038	0.122
Observations	76	76	76	76	76	76	76	76

4 Robustness Tests

4.1 Placebo Test

In this section we consider how unusual the Shale Discovery Announcement day is in generating the results we highlight in our main identification strategy. Given that the total stock market return is positive on this day, it is possible that the cross-section that we identify is really driven by market beta exposures (which are hard to measure). So we proceed by asking the question: How likely is it that a randomly picked trading day could produce a similar result?

To answer this, we repeat our analysis using all of the trading days in our post-crisis and shale oil period samples in place of the Shale Discovery Announcement day. We use the results to construct a “placebo” distribution of the key test statistics of interest. This distribution then allows us to assess how likely it is that a randomly-picked day would generate a portfolio that performs at least as well as the Shale Discovery Portfolio on the key economic dimensions that are most relevant to our study: the ability of the portfolio to explain the time-series of the U.S. aggregate stock market return during the shale period, as well as the relationship to the shale oil development itself. Specifically, we consider five statistics of interest. The first statistic measures the overall return on the total market portfolio during the shale oil period which is attributable to the shale discovery portfolio. (From Table 4, the relevant number is Change in Intercept, $\Delta\gamma_0^{Mkt} = 0.116\%$ in the last row of Panel B Column (8)). The second is the increase in the ability of the portfolio to explain variation in the market return, which is measured as the difference in the t-statistics of the slope coefficient of the shale discovery portfolio in explaining the market return γ_1^{Mkt} during the shale oil period relative to the t-statistic on this coefficient over the post-crisis period. (Table 4, first row Panel B, the difference in the t-statistics between columns (8) and (6) is 4.00). The third measures the ability of the portfolio to explain specifically Shale Oil Index returns in the shale period. To create the relevant statistic we take the difference in the t-statistics associated with $\gamma_1^{ShaleOil}$ and $\gamma_1^{S\&PInt}$. (Table 3, the t-statistic of the first row in column (8) of Panel D less the corresponding t-ratio in Panel B). The fourth statistic measures the ability of the portfolio to explain real shale drilling growth, and is the t-statistic in the regression of monthly growth

in the shale oil rig count on the Shale Discovery Portfolio return δ_2 , as shown in Table 2. The fifth is the t-statistic of the cross-section of industry returns in explaining employment growth, λ_E^1 (Table 5, first row of column (8)).

Figure 3 plots the histograms of these test statistics simulated using all of the trading days in both the post-crisis and shale oil periods (07/2009 - 09/2015). Panel A shows that the distribution of the market return explanatory coefficients is centered around zero but with a long right tail. Still, the Shale Discovery Announcement day is in the 95.0th percentile of this distribution. Panel B shows that the distribution of differences in the market explanatory power as measured as the difference t-statistics between the post-crisis (pre-shale) and the shale period is also centered near zero, but with a larger left tail. Consequently, there are fewer days that are “as good” or better at explaining this difference, so that the Shale Discovery Announcement is at the 98.6 percentile of this distribution. Panel C displays the distribution of the difference in t-statistics between the Shale Oil Index and the S&P Integrated Producer Index regression coefficient on the Shale Discovery Portfolio. These are also centered around zero, with the 3.9 value for the actual Shale Discovery Announcement day is at the 98.9 percentile of this distribution. Panel D plots the histogram of the t-statistics in the regression of monthly shale rig count growth on the lagged Shale Discovery Portfolio return. This distribution is centered around zero and quite dispersed. The Shale announcement day’s ability to link stock return to actual activity falls in the 96.3 percentile of all days. Panel E shows that the ability of the Shale Discovery day to explain employment growth is in the 95th percentile of all days. Finally, panel F summarizes the joint distribution of the test statistics described above by plotting the histogram of the lowest percentile of the five statistics corresponding to each day in the sample. Thus, it evaluates every day on its ability to explain the stock market returns jointly with the shale index returns, the shale drilling activity, and employment growth. By this measure, there is no other day in the sample that is as good as the Shale Discovery Announcement day that we use, whose minimum percentile (among the five values above) of 95.0 is the 100th percentile (highest of 1565 days) of the distribution. In fact, the second closest day has the minimum percentile of 87, which means that there is no other day in our sample that would fall in the 95th (or even 90th) percentile on all five measures.

This placebo evidence confirms that the shale discovery announcement is indeed a unique event, in that the industry stock returns on the day following this announcement have an unmatched ability to explain both the aggregate stock market return and employment growth during the shale period and to capture shale oil related news, both in the returns on shale stocks and in the real drilling activity.

4.2 Industry Market Betas and Shale Period Returns

In this section we reexamine the returns during the shale period from the perspective of market betas. We show that industry market betas, estimated either prior to the shale period or during the shale period, do not explain the large positive returns to the market post 2012.

The primary argument put forward in this paper is that the positive returns to the aggregate market post-2012 were driven by technological innovations in shale oil. Industries exposed to this shock experienced positive returns, while at the same time becoming systematically important to the market as a whole. For this reason, traditional “high beta” industries did not experience positive returns over this period.

As an illustration of this we perform a simple exercise. We construct characteristic portfolios using the cross-section of market betas estimated in each of the four sub-periods, and examine cumulative returns to this portfolio over the sample. Figure 4 plots the results. As the figure shows, these portfolios track the performance of the market very closely in the pre-crisis period, even those constructed using betas calculated in the latter half of the sample. More interestingly, the portfolios also track the market return very closely during the crisis and post-crisis recovery, but all subsequently exhibit a large divergence from the market beginning in 2012, consistent with the hypothesis that a new source of variation was driving market returns.¹⁴

¹⁴Unreported regression results show that all four portfolios exhibit high weekly correlation with the market return in all periods, but as all have negative or small positive returns in the shale period, none can explain a significant portion of the market increase over this period.

Figure 3: Placebo Tests for the Shale Discovery Announcement day

This figure shows the result of a placebo exercise in which other trading days' returns are used in place of the Shale Discovery Announcement day (8/1/2013). Each panel shows the histogram of a specific test statistic. To generate these histograms, we first obtain the cross-section of industry returns on an alternate trading day, and then repeat the analysis using these returns in the place of the shale discovery returns to recalculate the statistic. This exercise is repeated for each of the 1,595 trading days in the post-crisis and shale oil periods to create the distributions, which are plotted in blue. The red lines show the statistic obtained with the Shale Discovery Announcement day along with its associated percentile in the distribution of trading day statistics. The statistic examined in Panel A is the annualized return to the market portfolio return during the Shale Oil Period explained by the "trading day portfolios", calculated in analogous manner to the Shale Discovery Portfolio (The value for the Shale Discovery Portfolio is shown in Table 4). Panel B examines the difference in the t-statistics on the slope coefficient of the trading day portfolio return in explaining aggregate market returns during the shale period relative to the t-statistic on this coefficient in the post-crisis period (Table 4). Panel C shows the difference in the t-statistics of the slope coefficient on the trading day portfolio in the regressions explaining Shale Index returns and those explaining S&P Integrated Oil and Gas Producers index returns (Table 3). Panel D shows the t-statistic in the regression of monthly growth in the shale oil rig count on the trading day portfolio (Table 2). Panel E shows the t-statistic when using trading day returns to explain the cross-section of industry employment growth (Table 5). Panel F shows the minimum percentile for each trading day's place across the distributions shown in Panels A - E.


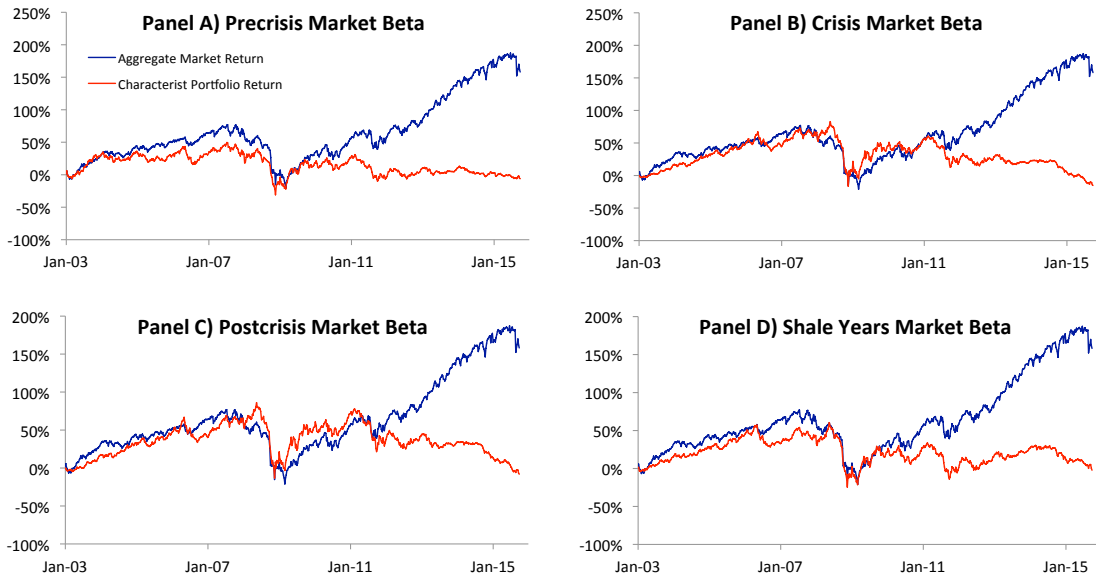


Figure 4: Market Betas and Cumulative Returns

This figure plots the cumulative aggregate stock market return against the cumulative return to characteristic portfolios constructed using industry market betas from the four sub-periods. The return on the characteristic portfolio in each week is the slope from a univariate Fama-Macbeth regression of that week's industry returns on a constant and each industry's market beta. This exercise is repeated four times with market betas calculated over the pre-crisis period (01/2003 - 06/2008), crisis period (07/2008-06/2009), post-crisis period (07/2009-12/2011), and the shale period (01/2012-10/2015), in Panels A), B), C), and D) respectively.



4.3 Monetary Policy Announcements

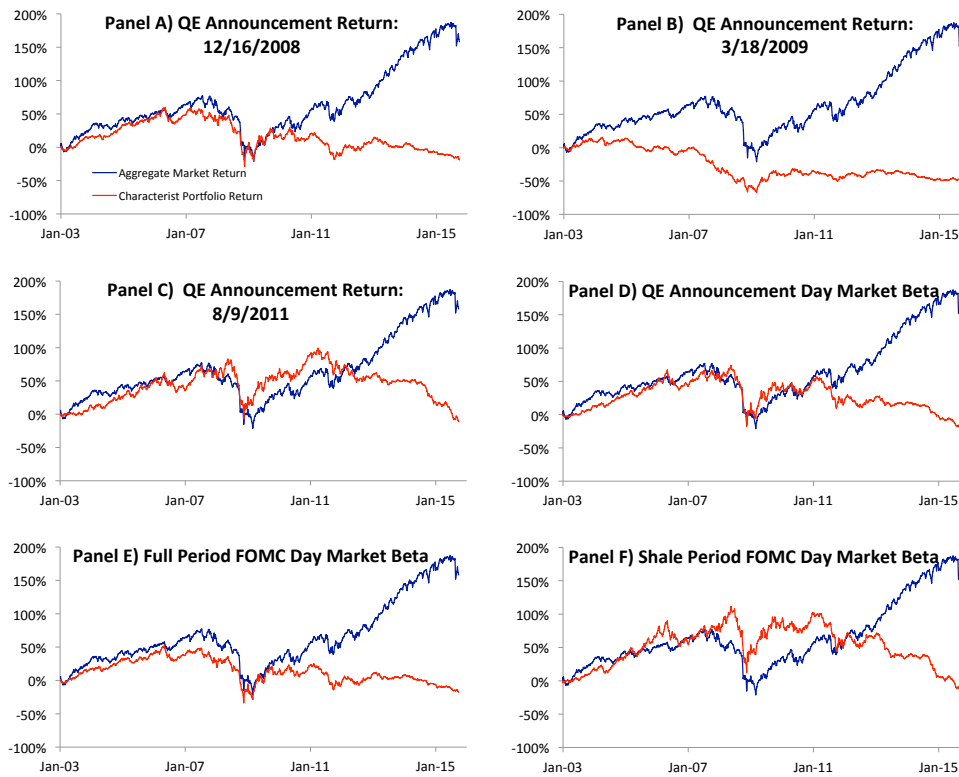
The period studied in this paper is noteworthy not only for technological innovations in oil production, but also for the implementation of unconventional monetary policy, such as the Federal Reserve’s Quantitative Easing program. While several recent papers examine the impact monetary policy on asset prices, to our knowledge this work has primarily utilized event studies focused on announcement day reactions (e.g Rogers, Scotti and Wright (2014) and Krishnamurthy and Vissing-Jorgensen (2011)), which shed little light on the cumulative impact of monetary policy over this period. In contrast, our methodology of examining characteristic portfolios formed using announcement day returns can be used to quantify the total contribution of monetary policy to aggregate stock returns. In this section, we use the cross-section of industry returns on monetary policy announcements to create characteristic portfolios, and find little evidence that monetary policy shocks can explain aggregate stock market increases in either the post-crisis or shale oil periods.

Including industry reactions to monetary policy announcements during the shale oil period also provides a test for robustness to the possibility that the results regarding our Shale Discovery Portfolio are driven simply by changes in the market beta of the relevant industries during this period. Savor and Wilson (2014) show that market beta is a good predictor of expected returns on stocks during days of monetary policy announcements by the Federal Open Market Committee, which are the days when the bulk of the equity risk premium is realized. As such, these FOMC announcement days are ideal for identifying potential non-shale shocks to U.S. stocks.

To proceed with these tests we first construct several industry-level proxies for exposure to monetary policy shocks using monetary policy announcement day returns. We examine FOMC announcements following Savor and Wilson (2014), as well as a group of 22 days regarding the federal reserve’s quantitative easing program following Wright (2012), Krishnamurthy and Vissing-Jorgensen (2011), and Neely et al. (2014). Using these days we construct six proxies. The first three proxies are the single day returns on the three quantitative easing days which saw the largest positive market return in the sample (12/16/2008, 3/18/2009, and 8/9/2011). The other three metrics are calculated as beta slopes β_j^{FOMC} from subsample time-series regressions of industry j returns on market returns using a group of announcement

Figure 5: Monetary Policy Exposure Measures and Cumulative Returns

This figure plots the cumulative aggregate stock market return against the cumulative return to characteristic portfolios constructed using various measures of industry exposure. The return on the characteristic portfolio in each week is the slope from a Fama-Macbeth regression of that week's industry returns on a constant and each industry's measure of Monetary Policy exposure. In Panels A), B), and C) industry monetary policy exposure is calculated as the industry return on three major announcements about the Federal Reserve's Quantitative easing program. In Panel D) the measure of exposure is the industry market beta calculated on 22 QE announcement days in the sample. In Panel E) the measure is the industry market beta calculated on all 95 FOMC meeting and announcement days in the sample. In Panel F) the measure is the industry market beta calculated on the subset of 12 FOMC days in the Shale Period.



days:

$$R_{t+1}^j = \alpha_j^{FOMC} + \beta_j^{FOMC} R_{t+1}^{Mkt} + \epsilon_{t+1}^{j,FOMC}. \quad (8)$$

Those groups are the 22 QE announcement days in the sample, the 95 days of FOMC announcements (both scheduled and unscheduled), and the 12 FOMC days which occur during the shale period.

As a first step we construct characteristic portfolios using univariate cross-sectional Fama-Macbeth regressions for each of our proxies. Figure 5 plots the cumulative returns to each of these portfolios. Each of these plots exhibits broadly the same pattern, with each characteristic portfolio resembling the pattern of the market beta portfolios shown in Figure 4. These portfolios generally track the market closely prior to the crisis, and then exhibit the same large drop and rapid recovery during the crisis period. However, all of the portfolios exhibit little increase in the post-crisis or shale periods, suggesting that monetary policy shocks have little explanatory power for the large increase in stock market valuations since 2010.

To show formally that these portfolios do not explain either the market returns or our findings regarding the shale discovery portfolio, we perform Fama-MacBeth regressions of industry returns on the shale and OPEC Announcements, this time including our three beta-based measures of industry exposure to monetary policy announcements.

Table 6 presents the results of multivariate cross-sectional regressions in Panel A. It is clear that the estimated impact of the shale announcement returns is completely unaffected by any of three measures, as all of the coefficients are essentially the same and the various monetary policy betas have no significant impact on the cross-section of industry returns. In Panel B, we construct new sets of mimicking portfolios using the slopes from these regressions, and repeat our analysis of the time-series performance of the total stock market. Panel B of the table shows that the monetary policy beta portfolios are quite strongly correlated with the market return over the shale period. However, the inclusion of these controls if anything strengthens the effect of the Shale Discovery portfolio on the market return. Additionally, the last line of the table shows the change in intercept when the monetary policy portfolio is removed from the regression. In no period are any of the three portfolios able to explain a significant portion of the market return.

These results show that the covariation between the aggregate stock returns and the shale

innovations that we identify using the Shale Discovery Portfolio are not driven by monetary policy shocks. In addition, the fact that market betas on FOMC announcements days in the Shale Period have little impact suggest that our results are unlikely driven by shocks that are altogether outside the shale oil sector, providing further validation for our approach.

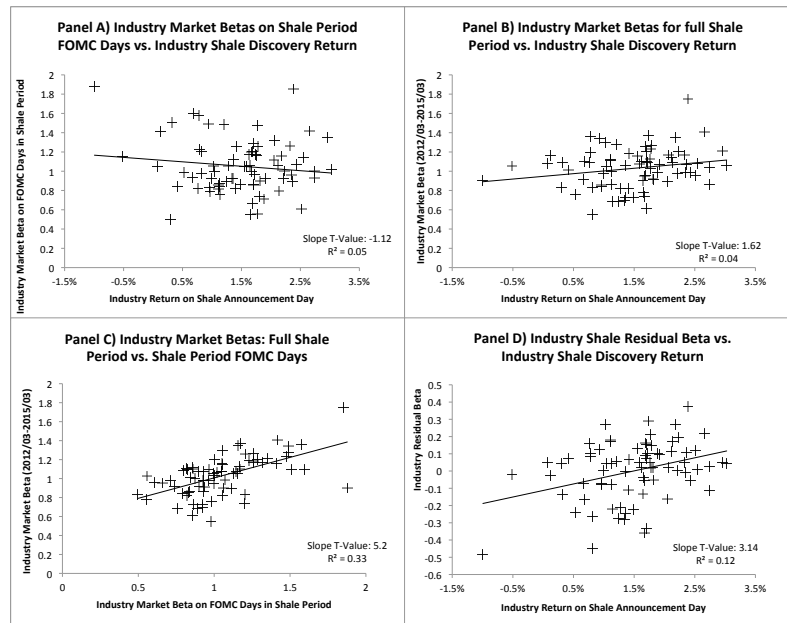
The exercise above is justified by the fact that the FOMC announcement day returns are indeed very closely related to industry market betas over the shale period. To illustrate this, Figure 6 provides plots of the relations between FOMC announcement day betas, market betas, and industry returns on the shale announcement day. Panel A plots the industry market beta calculated on the 12 FOMC meeting days in the Shale Period against the industry returns on the shale announcement day. As the plot shows, there is very little relation, again suggesting a distinction between the industries' response to the two type of shocks. Panel B then plots industry market beta calculated over the full shale period against the return on the shale announcement day. As the plot shows, there is a slight positive relation here, but it is not strong. Perhaps the most interesting plot is in Panel C, which shows that the FOMC announcement day beta is strongly related to the overall industry market beta over this period, with the former explaining 33 percent of the variation in the latter.

Finally, panel D shows that the shale announcement returns are able to explain a substantial of the variation in market betas not captured by the FOMC announcements (the plot shows the regression of residuals from panel C plotted against shale announcement returns). What is crucial for the validity of our identification is that the FOMC announcement returns do not line up with the shale announcement returns. If anything, they are negatively correlated, albeit weakly. Thus, it is not likely that the shale announcement returns are picking up some common macroeconomic shock that drives up asset prices over the shale period.

4.4 Alternate Measure of Shale Exposure

In this section we examine whether a simple and robust measure of shale industry exposure can yield similar explanatory power to our Shale Discovery Portfolio for the increase in the aggregate market value during the shale period. This test provides a verification that our results are not being driven by industries which are only tangentially related to shale oil (but happen to experience a high return on the Shale Announcement day).

Figure 6: Shale Announcement Returns, Market Betas, and FOMC Days



In this exercise, we first identify four groups of firms as the major shale-related industries. These groups are the Shale Oil Index, consisting of firms which own and operate shale oil leases; Drilling firms, which provide contract services to the primary oil firms; Oil Pipelines; and Railroads, both of which provide transportation of extracted shale oil from fields to refineries. We construct a Shale Dummy variable, which takes a value of one for these four industries, and zero for all others. We then repeat the primary analysis using this dummy as our metric for shale exposure. Table 7 shows the results.

The first four columns of Panel A show the results from cross-sectional Fama-Macbeth as in Panel A of Table 7, with the Shale Dummy used in place of the shale announcement return. What we see here is that the Shale Dummy firms outperformed the rest of the market in both the post-crisis and shale periods, after controlling for the OPEC Announcement returns, market beta in each of the first two periods, and the market beta on FOMC day in the shale period. In Panel B we repeat the exercise of Table 3 by constructing characteristic portfolios and examining their explanatory power for market returns. As the first four columns show, the shale dummy characteristic portfolio has a strong positive correlation with the market,

but only in the Shale Period. The last two rows of Panel B show the estimate of the change in intercept in the regression explaining the market return when the shale dummy characteristic portfolio is removed. Again we see that this estimate is positive and significant, suggesting that this portfolio has the ability to explain some of the positive market return but only in the shale period.

The last four columns of Panels A and B repeat this exercise, but this time include the shale announcement day return as a control the cross-sectional regression, and the corresponding characteristic portfolio is the time series regressions. While the shale dummy still exhibits significant explanatory power for the cross-section during the shale period, and the shale dummy characteristic portfolio is still positively correlated with the market over this period, the magnitude of both these coefficients is reduced. Most strikingly, the 0.084% per week explained by the shale dummy portfolio in column (4) is reduced to 0.025% per week in column (8). This suggests that the effect of the Shale Dummy is largely captured by the shale announcement return. These findings together imply that the two methods yield a similar result, and provide further evidence that the effects documented previously are in fact driven by the impacts of shale oil.

5 Conclusion

In a matter of a few years the technological innovations associated with fracking have revolutionized the U.S. oil market. We document that fracking innovations have had large effects on the aggregate stock market and the real economy. However, we caution that our estimates do not include or reflect the impact of potential costs or adverse consequences of shale development that are yet to be understood. Additionally, welfare consequences that are not linked with job growth or the stock market are not reflected in our estimates. Existing research has documented that shale oil and gas extraction has had localized negative impacts on the home values of houses dependent on ground water (Muehlenbachs, Spiller and Timmins (2015)). There have been additional unanticipated effects of shale development, such as an increases in seismic activity in shale producing states, such as Oklahoma. However, to the extent the negative environmental impacts are well understood already, our estimates

Table 7: An Alternative Proxy for Shale Oil Exposure

Panel A shows the results of Fama-Macbeth regressions of industry average returns. The first four columns shows regressions of sub-period average returns on a Shale Dummy variable, the industry return on the OPEC Announcement day (11/28/2014), the industry market beta calculated in the pre-crisis and crisis periods, and the market beta calculated on the 12 FOMC days during the shale period. The Shale Dummy variable takes a value of one for industries directly involved in the Shale Oil supply chain (Shale Oil Firms, Oil Pipeline firms, Drilling Service firms, and Railroads), and zero otherwise. Panel B shows the results of time-series regressions the market return on the characteristic portfolio returns constructed as the weekly cross-sectional slopes from the cross-sectional regression in Panel A. The final row shows the change in intercept in the time-series regression when the Shale Discovery characteristic portfolio is excluded.

Panel A: Fama-Macbeth Regressions of Industry Returns								
	Industry Average Returns				Industry Average Returns			
	Pre-Crisis (1)	Crisis (2)	Post-Crisis (3)	Shale Period (4)	Pre-Crisis (5)	Crisis (6)	Post-Crisis (7)	Shale Period (8)
Shale Industry Dummy	-0.129*	0.138	0.215**	0.321***	-0.097	0.238	0.167*	0.189**
	[-1.675]	[0.259]	[2.275]	[3.651]	[-1.310]	[0.533]	[1.773]	[2.326]
OPEC Announc. Returns	-0.192***	0.227	0.005	0.188***	-0.189***	0.237	0.000	0.175***
	[-3.415]	[0.608]	[0.074]	[3.340]	[-3.389]	[0.658]	[0.005]	[3.115]
Pre-Crisis Beta	0.210	-0.169	-0.007	-0.036	0.208	-0.174	-0.005	-0.030
	[1.316]	[-0.276]	[-0.038]	[-0.328]	[1.307]	[-0.286]	[-0.026]	[-0.270]
Crisis Beta	0.025	-0.291	0.022	-0.038	0.030	-0.275	0.015	-0.059
	[0.362]	[-0.279]	[0.124]	[-0.491]	[0.426]	[-0.264]	[0.082]	[-0.758]
Shale Years FOMC Day Beta	-0.104	0.372	-0.063	0.065	-0.099	0.389	-0.071	0.043
	[-1.115]	[0.854]	[-0.471]	[0.693]	[-1.048]	[0.898]	[-0.526]	[0.450]
Shale Discovery Return					-0.015	-0.047	0.022	0.062**
					[-0.785]	[-0.534]	[0.751]	[2.451]
Constant	0.139	-0.307	0.444***	0.298**	0.097	-0.055	0.386***	0.269**
	[1.317]	[-0.515]	[2.693]	[2.221]	[1.095]	[-0.099]	[2.862]	[2.264]
Observations	21,804	3,555	10,349	14,931	21,804	3,555	10,349	14,931
Number of Weeks	276	45	131	189	276	45	131	189

Panel B: Explaining Aggregate Market with Characteristic Portfolios								
	Aggregate Market Returns				Aggregate Market Returns			
	Pre-Crisis (1)	Crisis (2)	Post-Crisis (3)	Shale Period (4)	Pre-Crisis (5)	Crisis (6)	Post-Crisis (7)	Shale Period (8)
Shale Industry Dummy Portfolio	-0.278***	-0.163	-0.164	0.479***	-0.382***	-0.185	-0.247*	0.268**
	[-3.464]	[-0.696]	[-1.130]	[4.451]	[-4.636]	[-0.784]	[-1.657]	[2.371]
OPEC Announc. Portfolio	0.319***	0.075	0.061	-0.477***	0.365***	-0.024	0.074	-0.395**
	[2.789]	[0.214]	[0.264]	[-2.671]	[3.261]	[-0.064]	[0.326]	[-2.307]
Pre-crisis Beta Portfolio	0.591***	0.132	0.677***	0.460***	0.615***	0.163	0.697***	0.451***
	[19.769]	[1.086]	[10.039]	[6.296]	[20.711]	[1.278]	[10.349]	[6.465]
Crisis Beta Portfolio	0.527***	0.826***	0.893***	0.434***	0.585***	0.829***	0.883***	0.459***
	[7.814]	[13.627]	[13.004]	[3.963]	[8.708]	[13.596]	[12.991]	[4.395]
Shale FOMC Day Beta Portfolio	0.148***	0.046	0.193*	0.402***	0.164***	0.072	0.207**	0.374***
	[2.905]	[0.206]	[1.903]	[4.353]	[3.285]	[0.318]	[2.063]	[4.243]
Shale Discovery Portfolio					0.424*	0.740	0.582	2.194***
					[1.701]	[0.687]	[1.382]	[7.059]
Constant	0.065	-0.302	0.390***	0.224**	0.073	-0.290	0.389***	0.188*
	[0.895]	[-0.718]	[3.058]	[1.990]	[1.039]	[-0.685]	[3.082]	[1.749]
Weeks	276	46	131	189	276	46	131	189
R-squared	0.591	0.826	0.768	0.347	0.615	0.829	0.776	0.411
Market Return Explained by Shale Industry Dummy Portfolio								
Change in Intercept	-0.009	-0.011	-0.029	0.084**	-0.013	-0.015	-0.040	0.029
	[-0.589]	[-0.309]	[-1.088]	[2.124]	[-0.665]	[-0.399]	[-1.445]	[1.329]

T-Statistics in Brackets
*** p<0.01, ** p<0.05, * p<0.1

document that these effects are offset by positive employment and value effects in the stock market. The long run impact of shale technology on 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. In doing so we provide a novel framework for how to estimate the effect of technological innovations on the economy.

References

- Acemoglu, Daron, Amy Finkelstein, and Matthew J. Notowidigdo**, “Income and Health Spending: Evidence from Oil Price Shocks,” *The Review of Economics and Statistics*, October 2013, 95 (4), 1079–1095.
- Albuquerque, Rui, Martin S. Eichenbaum, and Sergio Rebelo**, “Valuation Risk and Asset Pricing,” *Journal of Finance*, *forthcoming*, December 2015.
- Allcott, Hunt and Daniel Keniston**, “Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America,” Working Paper 20508, National Bureau of Economic Research September 2014.
- Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng**, “News Shocks in Open Economies: Evidence from Giant Oil Discoveries,” Working Paper 20857, National Bureau of Economic Research January 2015.
- Back, Kerry, Nishad Kapadia, and Barbara Ostdiek**, “Slopes as Factors: Characteristic Pure Plays,” *Working paper, Rice University*, 2013.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis**, “What triggers stock market jumps?,” Technical Report, working paper 2015.

- 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.
- , **Dana Kiku, Ivan Shaliastovich, and Amir Yaron**, “Volatility, the macroeconomy, and asset prices,” *The Journal of Finance*, 2014, 69 (6), 2471–2511.
- Barsky, Robert and Lutz Kilian**, “Oil and the macroeconomy since the 1970s,” Technical Report, National Bureau of Economic Research 2004.
- Beaudry, Paul and Franck Portier**, “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, September 2006, 96 (4), 1293–1307.
- Bernanke, Ben S. and Kenneth N. Kuttner**, “What Explains the Stock Market’s Reaction to Federal Reserve Policy?,” *Journal of Finance*, 06 2005, 60 (3), 1221–1257.
- Blanchard, Olivier J. and Jordi Gali**, “The Macroeconomic Effects of Oil Shocks: Why are the 2000s So Different from the 1970s?,” Working Paper 13368, National Bureau of Economic Research September 2007.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen**, “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 2013, 81 (4), 1347–1393.
- 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.
- Campbell, John Y, Stefano Giglio, Christopher Polk, and Robert Turley**, “An intertemporal CAPM with stochastic volatility,” Technical Report, National Bureau of Economic Research 2016.
- Cascio, Elizabeth U. and Ayushi Narayan**, “Who Needs a Fracking Education? The Educational Response to Low-Skill Biased Technological Change,” Working Paper 21359, National Bureau of Economic Research July 2015.

- Cochrane, John H.**, “Presidential Address: Discount Rates,” *The Journal of Finance*, 2011, 66 (4), 1047–1108.
- Covert, Thomas R.**, “Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale,” 2014. Working paper, University of Chicago.
- Dvir, Eyal and Kenneth S. Rogoff**, “Three Epochs of Oil,” Working Paper 14927, National Bureau of Economic Research April 2009.
- Fama, Eugene F.**, *Foundations of Finance*, New York: Basic Books, 1976.
- and **James MacBeth**, “Risk, Return and Equilibrium: Empirical Tests,” *Journal of Political Economy*, May-June 1973, 81, 607–636.
- Feyrer, James, Erin T Mansur, and Bruce Sacerdote**, “Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution,” *Working Paper*, 2015.
- Gilje, Erik P.**, “Does Local Access to Finance Matter? Evidence from U.S. Oil and Natural Gas Shale Booms,” *Unpublished Working Paper*, 2011.
- Greenwald, Daniel L, Martin Lettau, and Sydney C Ludvigson**, “The origins of stock market fluctuations,” 2014. Working paper, National Bureau of Economic Research.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of Returns and Expected Returns,” *Review of Financial Studies*, 2014, 27 (3), 714–746.
- Gürkaynak, Refet S., Brian Sack, and Eric Swanson**, “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review*, March 2005, 95 (1), 425–436.
- 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.

- , “Causes and Consequences of the Oil Shock of 2007-08,” Working Paper 15002, National Bureau of Economic Research May 2009.
- 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.
- Hausman, Catherine and Ryan Kellogg**, “Welfare and Distributional Implications of Shale Gas,” *Brookings Papers on Economic Activity*, forthcoming, 2015.
- Jones, Charles M and Gautam Kaul**, “Oil and the stock markets,” *The Journal of Finance*, 1996, 51 (2), 463–491.
- Jones, Charles M., Owen Lamont, and Robin L. Lumsdaine**, “Macroeconomic news and bond market volatility,” *Journal of Financial Economics*, 1998, 47 (3), 315 – 337.
- 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 the 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.
- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, “The effects of quantitative easing on interest rates: channels and implications for policy,” Technical Report, National Bureau of Economic Research 2011.
- Lamont, Owen A.**, “Economic tracking portfolios,” *Journal of Econometrics*, November 2001, 105 (1), 161–184.

- Lamont, Owen and Andrea Frazzini**, “The Earnings Announcement Premium and Trading Volume,” Working Paper 13090, National Bureau of Economic Research May 2007.
- Lucca, David O. and Emanuel Moench**, “The Pre-FOMC Announcement Drift,” *The Journal of Finance*, 2015, 70 (1), 329–371.
- Muehlenbachs, Lucija, Elisheba Spiller, and Christopher Timmins**, “The housing market impacts of shale gas development,” *American Economic Review*, 2015, 105 (12), 3633–3659.
- Neely, Christopher J et al.**, “How Persistent Are Unconventional Monetary Policy Effects?,” *Federal Reserve Bank of St. Louis Working Paper Series*, 2014, (2014-004).
- Ready, Robert C**, “Oil prices and the stock market,” *Forthcoming: Review of Finance*, 2016.
- Rogers, John H, Chiara Scotti, and Jonathan H Wright**, “Evaluating asset-market effects of unconventional monetary policy: a multi-country review,” *Economic Policy*, 2014, 29 (80), 749–799.
- Sadorsky, Perry**, “Oil price shocks and stock market activity,” *Energy Economics*, October 1999, 21 (5), 449–469.
- Savor, Pavel and Mungo Wilson**, “How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements,” *Journal of Financial and Quantitative Analysis*, 4 2013, 48, 343–375.
- and — , “Asset pricing: A tale of two days,” *Journal of Financial Economics*, 2014.
- Savor, Pavel G and Mungo Ivor Wilson**, “Earnings announcements and systematic risk,” *Journal of Finance*, *forthcoming*, 2015.
- Wright, Jonathan H**, “What does monetary policy do to long-term interest rates at the zero lower bound?,” *The Economic Journal*, 2012, 122 (564), F447–F466.