

RFF REPORT

Oil Supply Shocks, US Gross Domestic Product, and the Oil Security Premium

Alan Krupnick, Richard Morgenstern, Nathan Balke, Stephen P.A. Brown, Ana María Herrera, and Shashank Mohan

NOVEMBER 2017



RESOURCES
FOR THE FUTURE

Oil Supply Shocks, US Gross Domestic Product, and the Oil Security Premium

Alan Krupnick, Richard Morgenstern, Nathan Balke, Stephen P.A. Brown,
Ana María Herrera, and Shashank Mohan*

Summary

The macroeconomic costs of unanticipated oil supply and oil price shocks remain the principal component of the oil security premium. A long history of academic papers have offered approaches to the estimation of such costs and the calculation of the oil security premium. Two relevant major changes have occurred in recent years: both the US economy and the world oil market are now more resilient, less dependent on oil in general, and (for the United States) less reliant on imports than a decade or two ago; and macroeconomic modeling has become more sophisticated, with advances coming from modeling dynamic economic relationships, using dynamic stochastic general equilibrium (DSGE) models, and extracting macroeconomic oil price shocks from time series data, using structural vector autoregression (SVAR) models. These advances suggest it is time to use sophisticated modeling tools to take another look at the macroeconomic effects of price shocks. In addition to using the DSGE and SVAR models, which are estimated directly from historical data, we also exercise the National Energy Modeling System (NEMS) model and perform a number of sensitivity analyses with all the models to check for the robustness of their estimates.

This report develops new estimates of the relationship among gross domestic product (GDP), oil supply and price shocks, and world oil demand and supply elasticities; translates them into oil security premiums using a welfare-theoretic-based computation model; and compares all these estimates with those in the literature. The literature is divided into three categories: older studies, newer ones, and a mixture of old and new.

* This effort was led by Alan Krupnick and Richard Morgenstern at Resources for the Future (RFF), who gathered an expert group of modelers and a distinguished set of advisors to help guide the effort. Our modeling partners included Ana María Herrera, University of Kentucky; Nathan Balke, Southern Methodist University; Steve Brown, University of Nevada, Las Vegas; and Shashank Mohan, Rhodium Group. Our advisors included Christiane Baumeister, Notre Dame; James Hamilton, University of California, San Diego; Martin Bodenstein, Federal Reserve Bank; James Stock, Harvard; and David Montgomery, RFF. We also would like to thank the Office of Energy Policy and Systems Analysis and the Office of Energy Efficiency and Renewable Energy of the US Department of Energy as well as the Sloan Foundation for their generous support of this project. Finally, we acknowledge helpful comments from a number of reviewers of this effort, especially Paul Leiby, Oak Ridge National Lab; Howard Gruenspecht (formerly of the EIA); and various other experts at the US Department of Energy, the US Energy Information Administration, and the US Environmental Protection Agency. The views and opinions of the authors expressed herein do not state or reflect those of the United States Government or any agency thereof.

© 2017 Resources for the Future (RFF). All rights reserved. No portion of this report may be reproduced without permission of the authors. Unless otherwise stated, interpretations and conclusions in RFF publications are those of the authors. RFF does not take institutional positions.

Resources for the Future (RFF) is an independent, nonprofit research institution in Washington, DC. RFF improves environmental, energy, and natural resource decisions through impartial economic research and policy engagement.

We find a wide range of estimates of the elasticity of GDP to an oil price change (the percentage change in GDP for a 1 percent change in the oil price) and the short-run price elasticities of supply and demand. Our new studies provide estimates of the former that overlap prior ranges of estimates but are generally lower in magnitude. Our new studies also find that world oil demand is more elastic in the short run than previously estimated. Overall, the implication is that the oil security premium is lower than that in the bulk of the existing literature. These values provide evidence that the changes in the economy are at least partly responsible for the lower values.

Yet a fundamental question remains. While we are able to connect the modern economy to the historical oil market, from the 1970s through what is termed the Great Moderation of oil price volatility (1984–present), we do not have real-world experience with major unanticipated oil market disruptions in these years. We use NEMS to make the connection between a 10 percent and a much larger (30 percent) price shock and find that the GDP elasticity and the resulting oil security premium are still lower than previous estimates. But there are questions about whether NEMS is sufficiently well suited for this task. Thus uncertainties remain about what an oil security premium would be in light of a large disruption. Advances in time and both theoretical and empirical modeling are needed to be more definitive about the macroeconomic effects of an oil disruption and its translation into an oil security premium.

Appendix

An appendix to this report includes the following sections and is available on the RFF website: <http://www.rff.org/research/publications/oil-supply-shocks-gross-domestic-product-and-oil-security-premium>.

- A. Oil Supply Shocks and the US Economy: An Estimated DSGE Model
- B. The Role of Oil Supply Shocks on US Economic Activity: What Have We Learned?
- C. Oil Price Shocks and the US Economy: An Application of the National Energy Modeling System
- D. New Estimates of the Security Costs of US Oil Consumption

Disclaimer

This report was prepared as an account of work sponsored by an agency of the US government. Neither the US government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the US government or any agency thereof. The views and opinions of the individual authors expressed herein do not necessarily state or reflect those of the US government or any agency thereof.

Contents

1. Introduction.....	1
2. What Is a Macroeconomic Effect of an Oil Supply Shock?	3
3. Brief History of the Literature on Oil Price Premiums.....	4
4. The Project Plan.....	5
Scenarios	6
Macroeconomic Models.....	6
Computational Model	6
5. Details of the Models.....	7
DSGE Model.....	7
SVAR Models	7
NEMS Model	8
Computational Model	9
6. Model Scenarios	10
7. Summary of Results	11
Elasticities	11
Oil Security Premium	15
8. Discussion.....	19
The Lack of Big Oil Supply Disruptions in the Modern Era	19
Asymmetry.....	21
Changes in the Oil-to-GDP Ratio	21
The Effects of Reduced Oil Imports	22
The Lack of Modeling to Address Foreign Oil Demand Shocks.....	22
Changes in the Short-Run Demand Elasticity for Oil.....	22
Conclusions.....	23
9. A Policy Perspective on the Oil Premium.....	23
References	25

1. Introduction

Beginning with the oil supply crises in the 1970s, the resulting long lines for gasoline, and the imposition of price controls, the US government and the public have focused on improving our energy security. The most visible improvement was probably the creation of the Strategic Petroleum Reserve. Other, more indirect changes were at least partly aimed at improving oil security, such as fuel economy standards for vehicles (which have environmental benefits as well). The analytic basis for all these policies was cost-benefit analyses of proposed government policies supported by serious academic and policy thinking about the value of increasing energy security—for what is not measured cannot credibly be accounted for.

Landsberg et al. (1979) introduced the idea that US dependence on imported oil would result in social costs that are greater than the market price paid for the oil. This thinking led to the identification of various benefits of greater oil security and then the assignment of economic values. These include the idea that oil supply disruptions, especially if unanticipated, can result in sizable losses in US real GDP. Potential components of the oil import premium include the macroeconomic risks associated with greater exposure to world oil supply disruptions, the effect of oil price shocks on transfers abroad, and a monopsony premium, which represents the US opportunity to exercise market power in buying oil on the world market.

The energy security literature suggests consumers are unlikely to consider how their own oil consumption affects the loss of economic activity resulting from world oil supply disruptions, which renders the expected losses in US real GDP as an externality. Brown and Huntington (2013) also identify the change in the expected transfers on the inframarginal barrels of imported oil associated with oil supply

disruptions as an externality. They argue that the expected transfers on the *marginal* barrel of imported oil occurring during a supply shock are something the purchaser can anticipate. But the change in the transfers on the *inframarginal* barrels of imported oil will not be anticipated. Accordingly, this report focuses on the macroeconomic and inframarginal transfer effects and does not address other benefits discussed in the literature, such as monopsony and military benefits.

For macroeconomic costs of unanticipated oil supply shocks, there is a long history of academic papers offering approaches to the estimation of such costs and the calculation of the oil security premium and to the estimation of the various ancillary relationships needed to make these calculations.


While these efforts have been ongoing, two relevant major changes have occurred in recent years. One is that both the US economy and the world oil market look very different today than they did even a few years ago. The other major change is that macroeconomic modeling has become more sophisticated. With respect to the first, the US economy is less dependent on oil than it was in the early 2000s, and with the fracking revolution and development of biofuels, our import share of oil consumption has fallen dramatically, dropping from 60 percent of US consumption to less than one-quarter today. At the same time, our economy is arguably more resilient now than it was a decade ago to shocks of any kind as a result of multiple factors, including increased global financial integration, greater flexibility of the US economy (especially labor and financial markets), reduced energy intensity of the US economy, increased experience with energy price shocks, and improved monetary policy. Collectively, this improved resilience of the US economy is

termed the Great Moderation (Figure 1).¹ On the world stage, OPEC's willingness and ability to create oil shortages has diminished, as the recent historical record of such supply-induced effects shows them to be much less severe than they were decades ago.

As for changes in macroeconomic modeling, advances have come from modeling dynamic economic relationships, using dynamic stochastic general equilibrium (DSGE) models, and extracting macroeconomic oil price shocks from time series data, using structural vector autoregression (SVAR) models. These advances suggest it is time to take another look at the macroeconomic effects of price shocks and to do so with sophisticated

modeling tools. The National Energy Modeling System (NEMS), a well-known simulation model, is also used in this project, as it is capable of translating very specific shocks into a wide array of results and can do so against explicit future projections of economic activity, using the *Annual Energy Outlook 2016* (AEO2016) as the baseline (EIA 2016). Sensitivity analyses are conducted to address various issues associated with using NEMS to model shocks. The results from these models are translated into oil security premiums using the Brown and Huntington computational model, which has undergone recent upgrading for new default values, particularly concerning the probabilities of oil supply shocks of various sizes.

FIGURE 1. REAL GDP AND RECESSIONS



¹ Note that even after the Great Recession of 2008-9, swings in the growth of GDP returned to a moderate trend.

Thus Resources for the Future (RFF) and its partners in academia (Ana María Herrera, University of Kentucky; Nathan Balke, Southern Methodist University; Steve Brown, University of Nevada, Las Vegas) and at the Rhodium Group (Shashank Mohan) have developed, with support from the US Department of Energy (DOE) Office of Energy Policy and Systems Analysis and Office of Energy Efficiency and Renewable Energy, new estimates of the macroeconomic impacts of oil supply shocks, the oil security premium, and some ancillary relationships needed to properly link the two.

This project began in March 2016. In addition to lining up the modeling team, RFF recruited a distinguished peer advisory group, including Christiane Baumeister (Notre Dame), James Hamilton (University of California, San Diego), Martin Bodenstein (Federal Reserve Board), James Stock (Harvard), and David Montgomery (RFF). The modeling team, the advisory group, and other invited guests from the government and academia participated in two face-to-face meetings, the first in July 2016 to review plans for the modeling and the second in December 2016 to review preliminary results. Substantial adjustments to the project were made in response to comments.

The rest of this report provides the necessary background for the reader to understand the issues and interpret the estimates, describes the modeling underlying the project, summarizes the results of the

analyses, and puts these results into perspective. Each modeler in the project wrote a detailed, stand-alone account of his or her methodology, including the structure, assumptions, estimation and simulation procedures, and data inputs used in the analysis. The modelers also presented their results for the GDP–oil price elasticities and other necessary elasticities, which are provided in full in the Appendix. The Appendix also presents the framework and results of the computational model, focusing on how the estimated elasticities translate into oil security premiums and how these premiums vary across modeling assumptions and over time.

2. What Is a Macroeconomic Effect of an Oil Supply Shock?

Mohan (see Appendix, Section C) lays out the channels for an oil supply change (or shock) to affect GDP. He argues that an abrupt or unexpected change in oil production (supply)² will raise oil prices and lead to reductions in oil consumption across all demand sectors: transportation, residential, commercial, industrial, and electric power.³ These shocks, even if ending quickly, can affect oil consumption both in the year when the shock occurs, by reducing utilization of current capital stock, and in future years, by diverting future capital stock toward more efficient equipment or fuel switching.⁴ For example, in the case of passenger cars, an increase in oil prices reduces total miles driven and induces shifts to more fuel-efficient or alternative-fueled vehicles, both of which reduce gasoline and diesel demands in

² An abrupt, unexpected increase in oil demand can also lead to increases in oil prices. However, we would expect supply-induced shocks to be contractionary and demand-induced shocks to be expansionary or at least less contractionary (see discussion of asymmetry below).

³ Because the electric power sector has foresight in NEMS, oil consumption changes even before the induced price shock. But since electric power accounts for less than 1 percent of oil consumption, this discrepancy is ignored.

⁴ The reverse of this will happen when there is a sudden increase in oil supply. In this study, we analyze the impact only of negative oil supply shocks—that is, sudden supply-induced increases in oil prices.

the near term. Buying more fuel-efficient or alternative-fueled vehicles in the near term changes the vehicle stock over time, which affects oil demand in the medium term.

Similarly, oil supply shocks affect GDP both during the impact year and in the future. In the impact year, higher oil prices push up nominal consumption expenditures—directly by increasing the price of liquid fuels and indirectly by increasing general inflation. This adversely affects real consumption expenditure. Higher oil prices also lead to higher investment in oil exploration and development and may crowd out investment in other sectors. International trade is also affected, as the relative prices of goods and services change as a result of the effect of high oil prices on the US economy and the economies of US trading partners. In response to the macroeconomic feedback from the oil price shock, the Federal Reserve may raise the federal funds rate,⁵ which affects borrowing costs and overall economic activity. Macroeconomic linkages carry these impacts forward to future years.

Brown and Balke (see Appendix, Section A) emphasize effects in the labor market. They argue that because capital and efficiency are fixed in the short run (or are subject to substantial adjustment costs in the medium term), a supply shock–induced increase in oil prices can affect output only through their effect on labor input. Thus the responsiveness of real GDP to induced oil price increases depends largely on the responsiveness of labor input and the elasticity of domestic oil supply. An increase in oil prices and the resulting decline in oil usage might cause a decline in labor demand, but the negative income effect

(given that the United States is a net importer of oil) would lead to an increase in labor supply. These two conflicting effects tend to mute the response of labor quantities and hence real GDP. Real wage rigidities would lead to larger changes in the quantity of labor, while adjustment costs in moving labor across sectors would dampen labor responses. Estimating the model over their sample period, Brown and Balke find that the overall response of hours worked to supply shocks is relatively small, suggesting substantial flexibility in how agents respond to oil supply shocks from the rest of the world (ROW). As a result, the real GDP response implied by the estimated model is relatively modest.

3. Brief History of the Literature on Oil Price Premiums

Since Landsberg et al. (1979) introduced the idea that US dependence on imported oil would result in social costs in excess of the market price paid for the oil, dubbing these costs the “import premium,” a number of others have reestimated this premium. These studies include EMF (1982), Bohi and Montgomery (1982a, 1982b), Broadman (1986), Bohi and Toman (1993), Parry and Darmstadter (2003), Toman (2003), and Leiby (2008). Some of the studies provided premiums under prevailing or projected world oil market conditions; others have estimated optimal oil import premiums that allowed market conditions to change in response to implementing the premium as a tax.⁶

The Council on Foreign Relations (2006) took a different approach and examined the political implications of US dependence on imported oil, identifying six different costs

⁵ Based on a built-in reaction function. See documentation of the Macroeconomic Activity Module (MAM) for further details at [http://www.eia.gov/outlooks/aeo/nems/documentation/macroeconomic/pdf/m065\(2014\).pdf](http://www.eia.gov/outlooks/aeo/nems/documentation/macroeconomic/pdf/m065(2014).pdf).

⁶ The optimal oil import premium would be lower than a premium estimated at prevailing market conditions because implementation of the tax reduces US oil imports and the world oil price.

associated with US dependence on imported oil. The study offers no guidance on quantifying these costs, however. In a departure from the previous economics literature, the National Research Council (NRC 2009) argues that the nonenvironmental externalities associated with US dependence on foreign oil over that of domestically produced oil are extremely small or nonexistent. The NRC defines what is meant by an externality and then proceeds to reject as externalities the macroeconomic risks associated with greater exposure to world oil supply disruptions, the effect of oil price shocks on transfers abroad, and the monopsony premium.

As noted above, Brown and Huntington (2013), partially following the National Research Council, identify the oil security premium as the macroeconomic losses and the expected transfers on the inframarginal barrels of imported oil associated with oil supply disruptions. In their computational analysis, they unsurprisingly find relatively smaller oil security premiums for domestic than for imported oil.

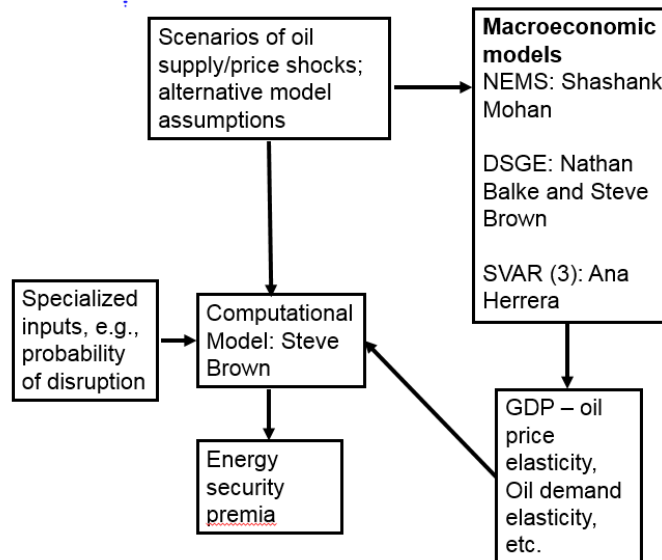
A recent report to Congress examines how energy security premiums are reflected in Regulatory Impact Analyses (RIAs) (DOE 2017). The report identifies a wide range of

federal actions that, in principle, relate to energy security. In practice, only a limited set of actions have addressed energy security through a quantitative analysis in RIAs, although a broader group of policies and actions have considered energy security in a qualitative way. The most complete and extensive discussion of energy security was done for the US Department of Transportation's and US Environmental Protection Agency's RIAs for rules setting fuel economy and greenhouse gas standards for cars and trucks, as well as for rules implementing a mandate for renewable fuels. That said, only benefits for consumers and the economy from reducing petroleum consumption (and oil imports) and environmental benefits from transitioning to a low-carbon economy are considered. There is only limited discussion of impacts on resilience, innovation, diversification, or other energy security goals. These benefits, taken from estimates of the oil security premium in Leiby (2008) for 2025, range from \$5 to \$8 per barrel.

4. The Project Plan

Figure 2 portrays how the project was organized and how the various pieces fit together.

FIGURE 2. ORGANIZATION OF THE PROJECT



Scenarios

The project logically begins with specification of scenarios of oil supply shocks or induced price shocks to be used in one way or another by all the models and a model-specific series of base case and alternative modeling runs to examine the sensitivity of outcomes to key parameter or assumption changes.

The most important elements of an oil shock scenario are the date, size, and duration of the oil supply shocks. The SVAR and DSGE models analyze historical supply shocks and their effects on GDP. NEMS requires assumptions about the effect of induced oil price shocks on GDP. Given the results of the empirical models, we set the severity of the shock at an initial 10 percent unexpected price increase in oil. This price increase was assumed to tail off over a 10-year period in NEMS following the price path implied by the DSGE modeling effort. Again, for NEMS, we assumed the shock would occur in 2030. This date was chosen to be far enough after the last year of applicable federal fuel economy and greenhouse gas vehicle standards to allow reasonable penetration of high efficiency vehicles.⁷ Thus in NEMS, the price shock is felt from 2030 to 2040 in runs that are compared with the AEO2016 reference case along with its price path. We also used NEMS to model a 30 percent induced price shock to look for nonlinearities in how the economy responds to larger price shocks.

The other scenarios are macro model specific. They are outlined below and each macro model is discussed in detail in the Appendix.

⁷ EPA has set final greenhouse gas vehicle standards for cars and light trucks out until model year 2025. (DOT has only issued an augural rule for the model year 2022–25 car and light truck standards.) DOT's and EPA's standards for medium and heavy-duty trucks extend out to model year 2027.

Macroeconomic Models

As noted, three alternative macroeconomic modeling approaches were used for this project: a DSGE model, an SVAR model, and the NEMS model developed by the Energy Information Administration (EIA) of DOE. These models generate or compute relationships between an oil supply–induced price shock and GDP in the form of an elasticity (the percentage change in GDP for a 1 percent change in the oil price), which is passed to the computational model. Other outputs from the three models are also passed to the computational model, depending on the information available from each macroeconomic model. These might include, for the United States, an oil demand and an oil supply elasticity, an income elasticity of oil demand, and demand and supply elasticities for the rest of the world

Computational Model

The computational model takes input from the other models to generate sets of oil security premiums, with each set reflecting one set of elasticities from one macroeconomic model (supplemented by default elasticities when the macroeconomic model is unable to generate a specific elasticity) and a time trend for the oil security premium from 2015 to 2040.

This model requires one other major set of inputs: the probabilities of an oil supply disruption of various sizes. These probabilities are taken from Beccue and Huntington (2016) based on a recent elicitation of experts. The probabilities are associated with supply disruptions ranging from 1 million to 21

million barrels per day (equivalent to 22 percent of current world oil consumption).

5. Details of the Models

DSGE Model

DSGE models, which rely on microeconomic principles such as specification of tastes, technology, and market structure, are estimated as systems based on historical data and designed to track the evolution of the economy over time and to handle shocks to the economy such as oil price changes and new technologies. They are general equilibrium models, in the sense that prices adjust to clear markets, but differ in that markets do not clear immediately, as adjustment costs for capital, habits in consumption, and labor are built in.

The particular model used in this project has several novel features. It adds real wage and nominal price rigidities, and it treats oil efficiency and intensity as endogenous. It also includes an endogenous domestic (US) oil supply, short- versus long-run demand elasticities, and endogenous adjustments in energy efficiency. The world oil price contains both endogenous and exogenous components that capture feedback from US economic activity and US oil production to world oil prices.

In most structural macro models that examine the interaction between oil prices and economic activity, oil affects the economy directly through consumption and capital services (which are a function of energy and installed capital). This model includes a transportation sector. In the nominal friction macro models, typically final goods are a composite good of many differentiated goods produced by monopolistic competitive firms. In this model, oil/energy is included in the production of the final good (in terms of producing transportation services), which is separate from oil in the production of the

differentiated good. This approach provides another margin through which oil will affect the economy. Private vehicles are included in the model through oil affecting consumption via their interaction with consumer durables. Key model parameters are estimated using Bayesian methods, which allow analysts to explore the sensitivity of estimates to alternative prior beliefs about these parameters. The model for this project was estimated using quarterly data for 1991 through 2015. This framework also allows characterization of uncertainty about the overall response of the US economy to oil price movements, appropriate elasticities of oil supply and demand, and sensitivity of this characterization to alternative priors. Once the parameters are estimated, the model can be used for simulation.

SVAR Models

Notwithstanding the name, these models contain much less structure than DSGE models in that estimation of a DSGE model (like many other macroeconomic models) requires many assumptions, whereas such assumptions are minimal for estimating an SVAR. The model involves regressing a vector of variables on their own and their lags, and it uses econometric techniques to sort out the complex relationships among variables rather than imposing a specific structure. It is designed to examine how shocks to an economy reverberate throughout the system and has previously been used to model oil price shocks.

In this project, three SVAR models for the global oil market are estimated using monthly data that span the period January 1973–December 2015 (see Appendix, Section B). The time series of structural oil supply innovations implied by each of the estimated models is extracted from monthly data and converted to quarterly time series. The quarterly time series of supply shocks is then

projected onto the log growth of US GDP, and impulse response functions are computed.

The three models are those of Kilian (2009), in which identification is attained through short-run restrictions; Kilian and Murphy (2012), in which both impact and dynamic sign restrictions are used for identification; and Baumeister and Hamilton (2015), in which more general prior beliefs are used to form priors on some parameters of the SVAR. A frequentist approach is employed to estimate the first two models, whereas Bayesian methods are used in the third. Note that these models differ not only in their identification assumptions but also on the basis of other modeling choices, such as the measure of global economic activity and the sample period covered in the original study. Hence, they provide a good testing ground for examining how different SVAR estimation strategies affect the size of the estimated response of US GDP to oil supply shocks.

NEMS Model

The models described above are designed to address dynamic forces in the economy set off by any shock to the economy. They are thus ideal for an analysis of the relationship between GDP and oil prices—the key elasticity in our work. Yet the SVAR model lacks almost any economic structure, and the DSGE model, though more complete in this respect than the SVAR model, is estimated, like SVAR, from historical data. Thus it cannot portray the modern economy or offer temporal specificity on the future timing and duration of a shock and subsequent effects on

GDP. For these effects, at the cost of losing the dynamic elements in these models and their tight grounding in historical experience, we need an energy simulation model.

The most frequently used and best-known energy model is NEMS, which is developed and maintained by the EIA. EIA primarily uses NEMS to produce the AEO, an annual publication that presents long-term projections of energy supply, demand, and prices in the United States.⁸

NEMS projects US energy production, consumption, and prices on an annual basis, subject to assumptions including but not limited to macroeconomic and financial factors, world energy markets, resource availability and costs, behavioral and technological choice criteria, and cost and performance characteristics of energy technologies. It is modular in nature, with each module of NEMS characterizing the future production, conversion, or consumption of energy in the United States. It uses a version of the Gauss-Seidel algorithm, where the model starts with a base solution and then iterates until it finds an equilibrium solution—a solution whose difference from the previous solution is less than a user-defined “tolerance” value.

NEMS is a structural energy model where relationships are based on historical data and empirical estimates and how they will evolve in the future. This is distinct from the DSGE and SVAR models, whose parameters are estimated as a system of equations from historical data. For instance, NEMS does not

⁸ The latest AEO, AEO2016, presents the forecasts through 2040. For more information on AEO2016 and earlier versions of AEOs, see <http://www.eia.gov/outlooks/aeo/>. NEMS source code is available to the public on request. The analysis presented in this report was performed using a version based on EIA’s source code and maintained by Rhodium Group.

estimate supply curves directly but, in effect, builds them up based on specific technologies and their costs.

The NEMS Macroeconomic Activity Module (MAM) provides both the macroeconomic and financial projections used in the model and incorporates the macroeconomic impact of changes in the energy system. It is divided into three submodules: US national economy, which provides national forecasts; industrial, which translates national forecasts into industry-level projections; and regional, which converts the results of the first two submodules into census-level forecasts. The US national economy module provides all the results used in this study. EIA uses a version of IHS's Global Insight (GI) model of the US economy to fill in this submodule. The GI model is an econometric dynamic equilibrium growth model. It incorporates insights from Keynesian, neoclassical, monetarist, supply-side, and rational expectations approaches. In addition, it includes the major properties of the long-term growth models presented by James Tobin, Robert Solow, Edmund Phelps, and others. This structure guarantees that short-run cyclical developments will converge to a robust long-run equilibrium. It includes the impact of interest rates and wealth effects on spending, thereby recognizing the importance of credit conditions on the business cycle and on the long-run growth prospects for the economy.⁹

NEMS is a fairly well-established model for analyzing US energy policy choices and market developments, and it uses a detailed representation of the US energy system for producing the forecasts. For example, to

forecast onshore US oil production, NEMS assesses the technical and economic constraints at the oil play level, which can be subregions or basins. Moreover, EIA uses the best available resources to update market and policy data and the structure of the model on an annual basis. Similarly, the MAM is updated by IHS to account for both near-term conditions and long-term structural changes in the economy and financial markets. Despite all these features, we recognize that NEMS is not as well suited to the task of modeling induced oil price shocks as the other models. As seen below, several steps were taken to increase our comfort level with the NEMS results, but this basic caveat still remains.

Computational Model

The computational model is a highly complex spreadsheet that relies on a welfare-theoretic model and is designed to combine results from the macro models with findings from the literature, including a schedule of the probability of future supply disruptions of various sizes (from Beccue and Huntington 2016) to estimate energy security premiums. The spreadsheet model used in this project was modified from the Brown and Huntington (2013, 2015) framework, which evaluates the expected costs of increasing the consumption of imported oil, increasing the consumption of domestically produced oil, or replacing domestic oil production with oil imports. The framework requires seven types of information: the elasticity of US GDP with respect to oil price shocks; the elasticity of non-US GDP with respect to oil price shocks; the probabilities of oil supply shocks of various sizes; the price elasticity of US oil demand (either by use category or in the

⁹ For full documentation of the MAM and other modules of NEMS, please refer to documentation shared at <http://www.eia.gov/outlooks/aeo/nems/documentation/index.cfm>.

aggregate); the price elasticity of non-US oil demand (in the aggregate); the price elasticity of US oil production; and the price elasticity of non-US oil production. Default values are included in the model but are replaced by values generated by the new macro model analyses wherever possible.

The computational model is parameterized for the AEO2016 reference case, an important choice that matches that of the NEMS modeling described above. It generates time-phased oil security premium estimates from the assumed date of the shock, presenting both a mean case and a probability distribution around that case, and does this for each scenario modeled.

6. Model Scenarios

Each modeling effort developed several scenarios for either directly modeling a given oil supply disruption or providing additional insights into drivers of such disruptions on GDP.

Mohan uses NEMS to evaluate the macroeconomic effects of six different sets of assumptions about oil supply disruptions. These include NEMS-RT, which is based on the AEO2016 reference case and a temporary disruption that leads to a one-period induced oil price shock; NEMS-RD, which is based on the reference scenario and an oil supply disruption whose effects on oil prices take place over 10 years; NEMS-RC, which is based on the reference case and a combination of a smaller response of non-US economic activity to induced oil price shocks, no US monetary policy response, and an oil supply disruption whose effects on oil prices take place over 10 years; NEMS-HC, which is based on EIA's high-price scenario and a combination of a smaller response of non-US economic activity to oil price shocks, no US monetary policy response, and an oil supply disruption whose effects on oil prices take place over 10 years; NEMS-LC, which is

based on EIA's low-price scenario and a combination of a smaller response of non-US economic activity to oil price shocks, no US monetary policy response, CAFE standards that are frozen at their MY 2017 level, and an oil supply disruption whose effects on oil prices take place over 10 years; and NEMS-RD30, which is identical to NEMS-RD except that the price shock in 2030 is three times larger (30 percent versus 10 percent).

These scenarios are chosen to tell a set of related stories about the size of the GDP effect (other things equal):

1. by duration of shock (NEMS-RT versus NEMS-RD)
2. by size of shock (NEMS-RD versus NEMS-RD30)
3. by oil price baseline time path (NEMS-RD versus NEMS-HC versus NEMS-LC)
4. by assumptions about ROW response to the shock and Federal Reserve response that would increase the GDP effect versus reference assumptions (NEMS-RC versus NEMS-RD)

For the set of elasticities identified as SVAR-BH, Herrera uses the Baumeister and Hamilton (2015) approach for oil price decomposition in a structured vector autoregressive model to estimate the effects of oil supply disruptions on US macroeconomic activity. Herrera also estimates models based on identifying assumptions and estimation procedures found in two other papers and examines the role of structural changes in the economy on the GDP elasticity.

Balke and Brown use two variations of a dynamic stochastic general equilibrium model to estimate the effects of oil supply disruptions on US macroeconomic activity. The set of elasticities identified as DSGE-S is based on standard preferences including a labor-leisure trade-off. Those identified as DSGE-GHH use Greenwood, Hercowitz, and Huffman preferences to exclude an income effect on

labor supply. The Balke and Brown DSGE model of the US economy also represents the world oil market, US international trade, and aggregate economic activity in the rest of the world (ROW). The model provides a mapping from structural shocks—such as those in technology, preferences, and oil supply—to observables such as oil prices, oil production, and other measures of economic activity. Balke and Brown use a combination of calibration and Bayesian methods to determine the model’s parameters and assess the stochastic process generating the exogenous shocks. The latter allows for the identification of exogenous oil supply shocks and the estimation of their effects on world oil prices and US real GDP. Balke and Brown conduct several additional simulations to test various hypotheses about the drivers of GDP elasticity.

7. Summary of Results

The results of this project are divided into two parts: elasticities used to compute the oil security premiums and the resulting values of the computed premiums.

Elasticities

Table 1 by Brown (replicated from the Appendix, Section D), shows the key elasticities used as inputs in the computational model, including the short-run price elasticities of world supply and demand, the US income elasticity of oil demand, and the elasticity of GDP with respect to a change in oil prices. Best estimates and upper and lower ranges are provided for all the elasticities, but statistical confidence intervals are provided at the 2.5th and 97.5th percentiles for SVAR-BH and DSGE-S. For simplicity, we drop results for DSGE-GHH and generalize results for the six NEMS scenarios into one, labeled NEMS. For details on the results from these additional scenarios, see the Appendix.

TABLE 1. PRICE, INCOME, AND GDP ELASTICITIES FROM THE INDIVIDUAL MODELS

Model	Short-Run Price Elasticity of World Supply	Short-Run Price Elasticity of World Demand	US Income Elasticity of Demand	Elasticity of US GDP with Respect to Oil Price Shocks
Benchmark-O	0.05 0.025 to 0.075	-0.055 -0.02 to -0.09	0.7 0.55 to 0.075	-0.044 -0.012 to -0.078
Benchmark-N	0.05 0.025 to 0.075	-0.175 -0.01 to -0.25	0.7 0.55 to 0.075	-0.018 -0.006 to -0.029
Benchmark-E	0.05 0.025 to 0.075	-0.055 -0.02 to -0.25	0.7 0.55 to 0.075	-0.028 -0.006 to -0.051
SVAR-BH	0.1526 0.0618 to 0.3162	-0.3554 -0.1797 to -0.7722		-0.0274 -0.0127 to -0.0623
DSGE-S	0.0582 0.0494 to 0.0736	-0.3328 -0.2808 to -0.4228		-0.007 -0.0064 to -0.0084
NEMS	0.2313 0.2129 to 0.2386	-0.2094 -0.2052 to -0.2123	0.8	-0.0197 -0.0128 to -0.0255

Sources: See Appendix.

Note that three benchmark sets of GDP elasticities have been developed by Brown from the extensive literature. Taken together, Benchmark-O (for old) and Benchmark-N (for new) yield good coverage of these estimated elasticities.¹⁰ Benchmark-O represents the older literature, which tends to offer higher oil security premiums that result from less elastic demand and to have a greater response of US GDP to world oil price shocks. Benchmark-N represents the newer literature (not counting the studies covered in the Appendix, see Sections A–C), with lower oil security premiums that result from more elastic demand and a lesser response of US GDP to world oil price shocks. Recognizing that these two sets of literature represent an *evolution* in thinking and modeling, but that the older literature has not been wholly overtaken by the new, Benchmark-E allows for a range of estimates to better capture the uncertainty involved in calculating the oil security premiums.

With the exception of the newer values of the short-run demand elasticities, the values in the table represent the Brown and Huntington (2013) interpretation of representative values from surveys by Atkins and Jazayeri (2004) and Dahl (2010a, 2010b) of oil demand elasticities, the Hickman et al. (1987) review of participating models in an Energy Modeling Forum study, the Jones et al. (2004) survey of the elasticities of GDP with respect to oil price shocks, as well as Krichene (2002), Cooper (2003), Huntington (2005),

Blanchard and Gali (2010), Leiby (2008), Hamilton (2009), Kilian (2009), Smith (2009), Balke et al. (2010), Kilian and Vigfusson (2011a), Kilian and Murphy (2014) and Baumeister and Hamilton (2015). The newer values of the short-run demand elasticity are the author’s adaptation of work by Davis and Kilian (2011), Kilian and Murphy (2014) and Coglianesi et al. (2015). The newer values of the elasticity of GDP with respect to oil price shocks are the author’s interpretation of work by Kilian (2009), Herrera and Pesavento (2009), Balke et al. (2010), Blanchard and Gali (2010), Kilian and Vigfusson (2011a and b), Kilian and Murphy (2014), and Baumeister and Hamilton (2015).

The size of the oil price shock originating from a supply disruption depends critically on the short-run oil demand and supply elasticities. More inelastic values of demand mean a greater price increase. As expected, the newer studies find that world oil demand is more elastic in the short run than previously estimated, with Brown’s central estimates from Benchmark-O to Benchmark-N more than tripling, from -0.055 to -0.175 . The values from the three macroeconomic modeling analyses conducted for this study are all higher than those revealed in Brown’s survey of the newer literature—especially those from the SVAR and DSGE models. The range of elasticities from these three modeling efforts is from -0.2052 to -0.7722 .¹¹

Our main results are for the elasticity of GDP with respect to induced oil price shocks.

¹⁰ Note that these new Brown baselines differ from the original Brown and Huntington (2015) estimates in a number of ways, as described in the Appendix (Section D). They are not US government estimates.

¹¹ As Brown notes in the Appendix (Section D), “Hamilton (2009) and Smith (2009) also provide compelling narratives about the movements in oil prices using very low elasticities of world oil demand. Consider Hamilton’s analysis of the 2004–2008 world oil market experience. Using the more elastic demand

values makes it impossible to track the path of world oil consumption with the actual prices and world GDP that prevailed at the time. In addition, consider the late-1973 oil supply disruption, which resulted in a 1.4 percent decrease in world crude oil supplies from the 1973 to 1974. World oil prices rose by 115.5 percent, which implies an elasticity of demand of -0.012 , and a more inelastic value if you consider the contraction in world economic activity.”

The older literature features a wide range of values (-0.012 to -0.078), with a point estimate of -0.044 (as summarized by Brown). The more recent empirical research—such as Kilian (2009), Herrera and Pesavento (2009), Balke et al. (2010), Blanchard and Gali (2010), Kilian and Vigfusson (2011a and b), Kilian and Murphy (2014), and Baumeister and Hamilton (2015)—suggests elasticities are likely to be at the lower end of the ranges used by Leiby (2008) (a central value of -0.032 , with a range of -0.01 to -0.054) and by Brown and Huntington (2015), with a point estimate less than half that of the older literature, in particular -0.018 (as summarized by Brown).

The corresponding elasticities from the three modeling analyses conducted for this study are roughly in line with those identified as the new literature (Benchmark-N), with central estimates ranging from -0.007 to -0.0274 . As noted, the differences between the newer and the older estimates may owe to improved monetary policy; increased familiarity with oil supply disruptions, making it easier for the economy to adjust; improved modeling techniques; and the lack of major oil supply disruptions in the past decade. Huntington (2016) cautions that the world has not seen a major oil supply disruption since 2003, which raises the possibility that research focusing strictly on relatively recent data is likely to give considerable weight to an era in which the phenomenon being studied has not occurred.

Some additional perspective can be gained by considering the results from the DSGE model. The elasticity of US real GDP with respect to an induced oil price shock is -0.007 (with a tight CI of -0.0064 to -0.0084), which is at the lower end of estimates in the literature. These estimates are fairly robust to changes in the model's specification. Using the estimated model in simulation mode, the authors examine a scenario where a negative

shock to ROW supply is scaled so that the real oil price rises by 10 percent during the first year after the shock. They find that following a shock, oil prices rise, peaking about two quarters after the shock, and then slowly return to the preshock level. US real GDP falls in response to the shock, with the peak decline occurring around the fourth quarter. Not surprisingly, ROW oil output falls in response to the supply shock, but the response is humped in shape, given the inertia estimated in ROW supply. At the same time, the US oil supply rises in response to increases in oil prices brought about by the decline in ROW supply. This response peaks around seven quarters after the shock and is relatively small, with an implied US oil supply elasticity in the first year after the shock of around 0.02. The model is also developed for a case where the vector shocks are augmented with a persistent ROW oil supply shock. Estimates of the posterior mode of the GDP/oil price elasticity (due to transitory shocks) for this model are virtually unchanged. This time path of oil price response for a persistent shock is provided to the NEMS model for its simulation for NEMS-RD and other related model runs.

Using the estimated DSGE model to conduct counterfactual analysis, Balke and Brown show that reducing the share of US oil imports below recent historical averages can substantially reduce the real GDP/oil price elasticity. This finding is important because it provides some evidence that changes in the US economy (increasing oil production and associated lower reliance on imports) are responsible for at least some of the reduction in the GDP elasticity.

The DSGE model is also exercised with a changed assumption that can be expected to boost the GDP elasticity—that is, that there is no income effect for leisure (Greenwood et al. 1988). The result is an estimated GDP/oil price elasticity that is only slightly higher than

in Balke and Brown's reference case. A version of the baseline model where the macro parameters are set equal to the modes of their prior distributions is also estimated. In this case, the mode of the GDP/oil price elasticity is estimated to be substantially lower than for the reference case.

Further, the model is used to examine whether the GDP response to a ROW oil demand shock is greater or less than the GDP response to a ROW supply shock. The GDP elasticity associated with the oil demand shock is slightly smaller than that for the supply shock.

Turning to the SVAR model results, in addition to the reference case (SVAR-BH), some variations were run to obtain additional insights. The major variation was in the SVAR model identification assumptions, which in turn imply different short-run elasticities of oil supply and demand. The reference case results are compared with those estimated in Kilian (2009) and Killian and Murphy (2012). Using identical data, these last two sets of assumptions and estimation procedures generate even lower GDP elasticities in response to oil supply disruptions than those from the reference case. This is due to larger short-run elasticities of oil demand and lower elasticities of supply. The dynamic response of real GDP also differs across specifications—a slower and longer-lasting impact for the reference model and an immediate but sharply diminishing impact for the two alternative sets of assumptions and estimation procedures. That is, Herrera finds that specifications where the short-run elasticity of oil supply is assumed to be very close to zero and the elasticity of demand is larger result in a smaller and

shorter-lived negative effect of oil supply disruptions on US GDP.

Additional results from the SVAR model suggest that structural changes to the US economy have contributed to the lower GDP elasticities. Herrera (see Appendix, Section B) notes a literature (Blanchard and Gali 2010; Edelstein and Kilian 2009; Herrera and Pesavento 2009; and Herrera and Karaki 2015) showing that induced oil price shocks are having a more muted effect on GDP since what is termed the Great Moderation, the reduction in the volatility of business cycle fluctuations starting in the mid-1980s and appearing to return after the Great Recession of 2008–9. This more recent period coincides with a decline in the volatility of crude oil prices and a reduction in the share of energy in personal consumption expenditures. Herrera then uses the model in two ways to address this question: by dropping periods earlier than 1984, when volatility and shocks were greater than in recent years; and by performing a recursive analysis that starts with data from 1975 to 1990 and reruns the model each time an additional quarter of data is added. She finds that the GDP elasticities are much lower as one adds more recent periods or examines the more recent period relative to the entire period. As these regression-based estimates hold other things constant, the implication is that the economy's structure is driving the observed reductions in GDP elasticity.

The results for the six NEMS runs are also instructive. NEMS-RD provides a GDP elasticity larger than NEMS-RT (-0.0255 versus -0.0195)¹² because the former imposes a 10-year (declining) duration for the shock versus the latter's one-year duration. Interestingly, the difference is not large.

¹² Here we mean a larger negative but dispense with the full term for conciseness.

The GDP elasticities applying to different sizes of shocks are important because shocks have been small during the Great Moderation period. So the NEMS-RD versus NEMS-RD30 elasticities can provide some information on whether the GDP effects scale proportionally, or less or more than proportionally, to the size of the price change. The answer is less than proportionally (-0.0255 versus -0.0208).¹³ Recall, however, the caveat that NEMS is not well suited to modeling these induced price shocks.

In various respects, the NEMS model responds as anticipated. The effect of changing ROW oil supply elasticities and assuming no Federal Reserve response to the induced oil price shock would be expected to lower the GDP elasticity, and in fact, it does (-0.0177 versus -0.0255). A smaller response of non-US economic activity to induced oil price shocks means higher net exports from the United States, while no US monetary policy response to increased inflation means lower interest rates. Both of these effects lower the impact on GDP and hence reduce the GDP elasticity.

The effect of different price paths is as expected. The hypothesis is that a higher price path, given a percentage shock, would translate into a larger absolute shock than for the reference case (-0.0231 versus -0.0177). Conversely, with a low oil price path assumed,

the GDP elasticity related to the low oil price baseline is considerably lower than its reference case counterpart (-0.0128 versus -0.0177).

Oil Security Premium

Figure 3, which draws on work by Brown (see Appendix, Section D), provides the average value of the aggregate oil premiums over the 2015–40 time horizon for the benchmark scenarios and the three new modeling efforts. As discussed, these oil security premiums are based strictly on well-specified externalities and include only the change in the expected GDP loss from an additional barrel of oil consumption plus the change in the expected transfers on the inframarginal barrels of imported oil. As shown in Figure 3 and Table 2, none of the models yields results close to Benchmark-O, and the new model results detailed in the Appendix yield oil security premiums below those of the benchmarks. As shown by Brown (see Appendix, Section D), the change in expected GDP loss for a marginal change in imported oil consumption is vastly larger (in \$/barrel terms) than the change in expected transfers on inframarginal oil imports for instance, for Benchmark N (mid), the GDP loss is \$1.2 per barrel and the inframarginal transfer is \$0.006 per barrel.

¹³ Shashank Mohan explains this result as follows (see Appendix, Section C): Oil expenditures do not proportionally increase when the induced price shock increases from 10 to 30 percent because oil demand falls. This makes the impact on real consumption smaller than one would expect if oil expenditures grew in line with price changes. Moreover, as the short-term domestic supply elasticity is lower for the 30 percent case, the investment in oil exploration and development also exhibits slower growth with the change in oil prices, leading to a proportionally smaller impact on

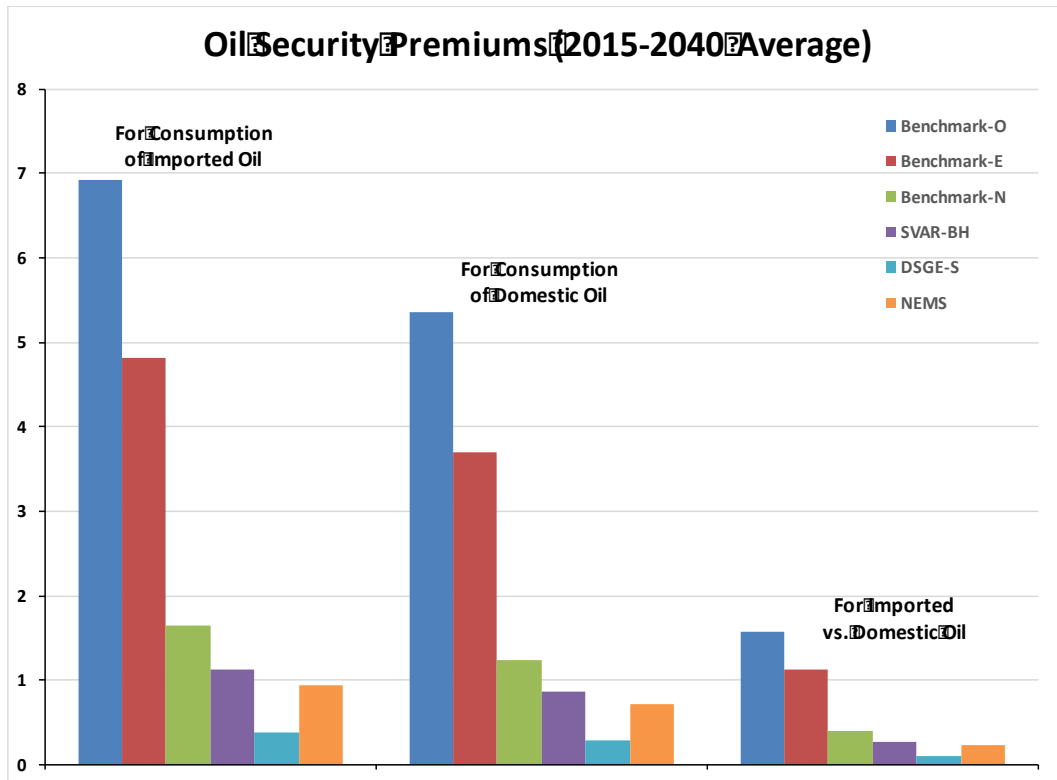
investment. Smaller further changes in consumption and investments lead to a smaller GDP elasticity under the 30 percent shock than under the 10 percent shock. Net exports also grew more slowly, which would lead to an increase in GDP elasticity, ceteris paribus. But that was not enough to compensate for smaller changes in consumption and investment, and on net, the GDP elasticity is smaller under a 30 percent shock scenario than under a 10 percent shock scenario.

TABLE 2. AGGREGATE OIL SECURITY PREMIUMS, 2015–40 AVERAGE (US\$2015 PER BARREL)

Model	Consumption of Imported Oil	Consumption of Domestic Oil	Imported vs. Domestic Oil
Benchmark-O	\$6.92 \$1.47 to \$20.03	\$5.36 \$1.10 to \$15.73	\$1.56 \$0.37 to \$4.30
Benchmark-N	\$1.64 \$0.77 to \$4.50	\$1.25 \$0.58 to \$3.46	\$0.39 \$0.19 to \$1.04
Benchmark-E	\$4.83 \$0.40 to \$15.62	\$3.70 \$0.29 to \$12.21	\$1.13 \$0.11 to \$3.41
SVAR-BH	\$1.12 \$0.25 to \$4.84	\$0.86 \$0.19 to \$3.76	\$0.26 \$0.06 to \$1.08
DSGE-S	\$0.39 \$0.28 to \$0.54	\$0.28 \$0.20 to \$0.40	\$0.11 \$0.08 to \$0.14
NEMS	\$0.94 \$0.60 to \$1.27	\$0.72 \$0.46 to \$0.97	\$0.22 \$0.15 to \$0.30

Source: Model estimates.

FIGURE 3. AGGREGATE OIL PREMIUMS FOR VARIOUS MODELS (2015–40 AVERAGE)



Figures 4–6, based on work by Brown (see Appendix, Section D), show how the aggregate oil premiums evolve over the 2015–40 time horizon. As shown in Figure 4, all the models show an increasing oil security premium for the consumption of imported oil. Gains in the change in the expected GDP loss from increased consumption of imported oil more than offset the change in the expected transfers for inframarginal oil imports resulting from increased consumption of imported oil. As shown in Figure 5, all the models show an increasing oil security

premium for the consumption of domestically produced oil. Gains in both the change in the expected GDP loss and the change in the expected transfers for inframarginal oil imports resulting from increased consumption of domestically produced oil account for the increase. As shown in Figure 6, the oil security premiums for the substitution of imported oil for domestically produced oil generally rise slightly for Benchmark-O and are generally constant for Benchmark-N, Benchmark-E, and the NEMS, SVAR, and DSGE models.

FIGURE 4. OIL SECURITY PREMIUMS: US CONSUMPTION OF IMPORTED OIL




FIGURE 5. OIL SECURITY PREMIUMS: US CONSUMPTION OF DOMESTIC OIL

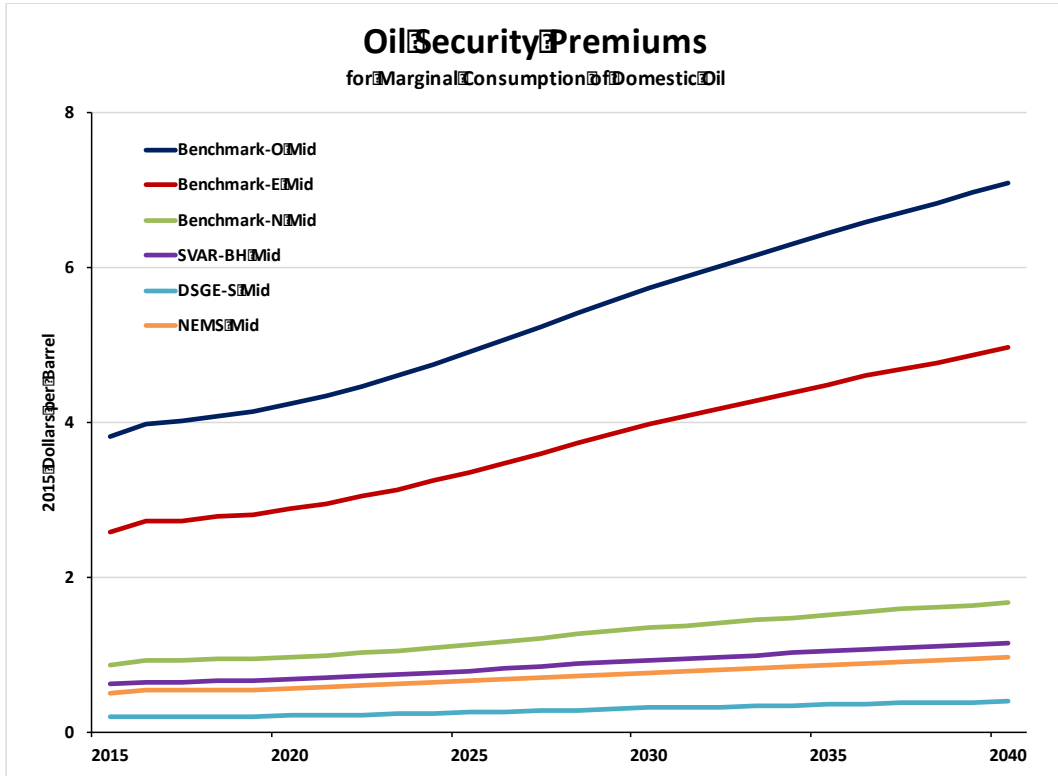



FIGURE 6. OIL SECURITY PREMIUMS: IMPORTED VERSUS DOMESTIC OIL



As demonstrated by the differences in the premium estimates for Benchmark-O and Benchmark-N, the elasticities from the newer economics literature suggest much lower premium estimates, with the Benchmark-E estimates in between, by design. The results from the NEMS, SVAR, and DSGE exercises are consistent with the newer literature. As noted, calculating oil security premiums with these three macroeconomic modeling exercises has the advantage of using sets of elasticities that have been jointly estimated with the state-of-the-art models that yield parameter values that are internally consistent within each model.

8. Discussion

Despite the wide range of estimates, we recognize that these calculations of the oil security premiums are substantially lower than the values used in RIAs conducted over the past decade, which for the most part have been computed from older literature. What accounts for the lower GDP/oil price elasticities, as well as the larger short-run demand responses that drive the calculations of these oil security premiums?

Why, the authors ask, is the GDP to oil price elasticity so low, especially relative to the earlier empirical literature? The general equilibrium approach taken by this report implies that all prices—not just the price of oil—respond when there is an oil supply disruption. The price responses throughout the model generally lower the magnitude of quantity responses (for variables such as nonoil goods and number of hours worked) compared with what would be the case if prices and wages did not change. The larger price responses and lower quantity responses tend to reduce the elasticity of real GDP with respect to real oil price changes. Despite some wage and price stickiness and various types of adjustment costs, the model still finds substantial flexibility for economic agents to

adjust to oil price changes. Thus less sticky wages and prices in the currently structured economy could account for lower GDP effects. Indeed, Blanchard and Gali (2010) argued that a declining oil-to-GDP ratio, increased labor market flexibility, and improved monetary policy have all contributed to the declining importance of oil price changes in macroeconomic fluctuations.

Below, we discuss six issues that help put these estimates into perspective: (1) Does it matter that the historically large oil supply disruptions are concentrated in an earlier period that is not well covered in the most current estimates of the short-run elasticities of demand or the elasticity of GDP with respect to oil prices? (2) Should one be concerned about the possibility of an asymmetric response of US real GDP to the direction of induced oil price shocks? Or, put another way, would an oil price decrease following a positive supply shock have the same effect (with opposite sign) on GDP (and the corresponding oil security premium) as an increase in the oil price? (3) Is the reduced US oil-to-GDP ratio responsible for the attenuated GDP response? (4) Have reduced US oil imports weakened the response of US real GDP to oil supply disruptions? (5) Is it important to address US exposure to foreign oil demand (as opposed to supply) shocks to fully assess the risks to US oil consumption? (6) How have short-run demand elasticities changed, and why does this matter?

The Lack of Big Oil Supply Disruptions in the Modern Era

A lower oil supply (and induced price) change paired with a more flexible, less-oil- and oil-import-dependent economy will logically lead to smaller effects on GDP and a lower oil security premium. On the one hand, as Huntington (2016) cautions, the world has not seen a major oil supply disruption since 2003, which raises the possibility that the

newer research is providing elasticity estimates that would not apply in the case of a major oil supply disruption. On the other hand, considering the differences between the modern and older economy, the current effects of any given oil supply disruption are probably now smaller than was estimated with data from the time period in which the big oil supply shocks occurred. Oil consumption has likely become more flexible. The economy has become better able to adjust to oil price shocks, consumers and businesses better know the effects of oil supply disruptions, and monetary policy is better informed about how to respond to supply disruptions.

We have examined the proposition that the economy has undergone structural change leading to lower GDP effects of a disruption. As noted above, Herrera (2016) uses both a recursive and a two-period analysis to show that the elasticity of GDP with respect to oil prices declined over the time period from the 1990s to the 2010s, holding the size of the oil shock constant. This result provides some evidence that the economy is less vulnerable to an oil supply shock.

Nonetheless, big supply disruptions may put more stress on economic relationships than the small oil supply disruptions we have seen in recent years. Consequently, the large oil supply disruptions, which are outside the estimation range of the models, might generate more inelastic supply and demand responses and a stronger GDP response to any given oil price shock than are found with the models using recent data. The result would be a greater price shock and a bigger GDP loss. The resulting oil security premiums would be larger, closer to the values estimated for Benchmark-O.

Interestingly, our modelers have raised an alternative hypothesis. Because economic adjustments are costly, adjusting to small oil shocks might not be worth it, whereas adjusting to bigger shocks would be. In this

case, the GDP elasticity for a big shock could be less than that for a small shock.

Given that we have not observed a modern economy with large oil supply disruptions, there is no reliable method to quantify what these differences might be on the basis of historical data. Nonlinear models might be used to evaluate how the elasticities change with the size of disruptions, but such an approach would not be based on any actual observations of large disruptions in a modern economy.

Extending the data used for estimation farther back in time creates a different problem. Estimation over a long time span increases the possibility of structural change that is not captured by the model. At best, the result would be an average of old and new results. At worst, the result would involve greater estimation challenges and be a poor fit.

If we postulate a world in which the economy responds to small oil supply disruptions in a manner that is well captured by the newer estimates and responds to big supply disruptions in a manner that is better captured by the older estimates, we can consider an exercise in which the elasticities used to evaluate the security premiums change with the size of the disruptions. We could use elasticities from the newer literature for small oil supply disruptions and elasticities more similar to that found in the older literature for the big oil supply disruptions, with graduated intermediate elasticities to cover the transition from small to big disruptions. Although, we do not know exactly how to set the transition between parameters, the resulting oil security premiums would inevitably lie somewhere between the smaller estimates found with Benchmark-N and the bigger estimates found with Benchmark-O. The estimate using elasticities that vary with the size of the disruption will show greater price and GDP effects for the bigger disruptions than are found with Benchmark-N. Similarly, the

elasticities will show smaller price and GDP effects for the smaller disruptions than found with Benchmark-O. Hence, Benchmark-N and Benchmark-O represent reasonable bounds by which the effects of oil supply disruptions might affect the economy.

This dilemma led us to use NEMS in the modeling exercise as a way of capturing some of the new economy characteristics, recognizing that the NEMS structure lacks the dynamic adjustments to shocks in other models. Thus the only direct evidence we have for whether the new economy might respond differently to a big shock than a small shock is from two NEMS runs: NEMS-RD and NEMS-RD30. NEMS-RD models a 10 percent oil price rise, and NEMS-RD30 models a price rise of 30 percent. If the resulting GDP impact of the latter were more than three times that of the former—in other words, if there is a non-linear response of GDP to price changes—then that would lend support to the view that our elasticities are too low. Actually, the GDP elasticity is lower for a 30 percent change than for a 10 percent change in oil price.

We note that use of the DSGE and SVAR models in this project to test nonlinearities is far from a trivial exercise and, in fact, would represent major advances in the literature. We would encounter two challenges. First, relaxing linearity assumptions involves significant computational challenges. Second, modeling exactly how nonlinearities would occur and how to represent them in the model involves challenges to theory development. Thus we have concluded that additional work is needed to better understand the potential for nonlinear responses of GDP to price changes.

Asymmetry

From the late 1980s until late in the first decade of 2000, a considerable body of research found that US GDP responded asymmetrically to world oil price shocks, with

increased prices having a bigger negative effect on economic activity than the positive effect from decreased prices. Contributions include Mork (1989), Hamilton (1996, 2003), Davis and Haltiwanger (2001), and Balke et al. (2002). Since Kilian and Vigfusson (2011a, 2011b) specified a new set of tests for asymmetry and macroeconomic modelers began using newer data sets, however, no peer-reviewed articles have found an asymmetric relationship between oil prices and US GDP.

Identifying asymmetry is an issue similar to evaluating the potential effects of bigger oil supply disruptions. In the older literature, which relied on older data sets, these asymmetric tests were important evidence for the finding that US GDP responded strongly to oil price shocks resulting from supply disruptions. In the newer literature, which is specified with symmetry and relies on data sets that mostly exclude big disruptions, the elasticity of GDP with respect to oil price shocks has been much lower. As noted, we have not seen large supply disruptions in our modern economy or in the newer analyses used to evaluate asymmetry. Perhaps an asymmetric response occurs only as a result of large supply disruptions.

Changes in the Oil-to-GDP Ratio

From 1973 to 2015, the US oil-consumption-to-GDP ratio fell by more than 60 percent. Has the decline in this ratio reduced the sensitivity of the US economy to oil price shocks originating from oil supply disruptions, as might be expected? The answer is unclear. As of yet, no one has produced an empirical paper showing that the reduced oil-consumption-to-GDP ratio for the United States has weakened the response of US real GDP to oil supply disruptions.

Can an answer be found by looking across various developed countries with different oil-to-GDP ratios? Not definitively. For countries

other than the United States, the published research on the response of real GDP to the price effects of world oil supply disruptions is more limited. The economies of most other countries that have been studied also are much more open than the US economy, and macro monetary policy has been conducted differently. And few studies have addressed multiple countries in a single analysis to ensure consistency of approach. Examining eight OECD countries, Brown et al. (1996) find preliminary evidence that oil-importing countries that had higher oil-to-GDP ratios also faced more difficult trade-offs in inflation and GDP losses in response to oil-price shocks than did oil-importing countries with lower oil-to-GDP ratios. With only preliminary evidence, however, the question remains quite open.

The Effects of Reduced Oil Imports

From 2005 to 2015, US dependence on oil imports declined from 60 percent of domestic consumption to 24 percent, and US reliance on oil imports is projected to decline further in the AEO2016 (and AEO2017) (reference case, although the projections in side cases span the space from the US being a significant net importer to significant net exporter, depending on price, resource, and technology assumptions. Do these declines in oil imports reduce the vulnerability of the US economy to world oil supply disruptions? The answer is probably yes. Reduced US oil imports have been the result of increased US oil production. These increases in US oil production increase the share of stable oil supplies in the world oil market and thereby cushion the price effects of a given disruption, an effect that is captured in the present analysis.

What reduced reliance on oil imports does not do, however, is prevent an induced oil price shock from reaching the United States. Because oil is a fungible commodity, the price shocks resulting from supply disruptions

elsewhere in the world are transmitted to the US economy without regard to the quantity of oil that is imported. As the United States moves toward zero net oil imports, however, the losses in the sectors of the economy that are hurt by induced oil price shocks will be increasingly offset by the gains in the sectors of the economy that benefit from oil price shocks. Brown and Yücel (1995, 2013) have quantified these effects at the state level, and Balke and Brown (2016) show that reducing the share of US oil imports below recent historical averages can substantially weaken the response of US real GDP to oil prices.

The Lack of Modeling to Address Foreign Oil Demand Shocks

Oil security premiums rely on estimates of the price effects of world oil supply disruptions but do not take into account probable foreign demand shocks. Is this an oversight in evaluating energy security? Identifying foreign oil demand shocks as an external security cost of oil consumption may be less important. Unexpected growth in global oil demand (perhaps driven by the business cycle) is not likely to be experienced as sudden oil price movements because nothing on the demand side changes quickly, except when there is a Fukushima-like event that shifts a country's electric power generation from nuclear power plants to those that are oil-fired. There also seems to be no reason to be more concerned about the effects of international business cycles affecting the US economy through variations in oil demand than through any other channel by which these effects are transmitted.

Changes in the Short-Run Demand Elasticity for Oil

As noted by Brown (see Appendix, Section D), the size of the oil price shock originating from a supply disruption depends critically on the short-run elasticities of demand and supply. More inelastic values of

demand mean a greater price increase. An older literature, including surveys by Atkins and Jazayeri (2004) and Dahl (2010a and 2010b), as well as narratives by Hamilton (2011) and Smith (2009), finds that short-run oil demand is very inelastic, with Brown and Huntington (2013, 2015) using -0.055 in a range of -0.02 to -0.09 . More recent econometric studies of oil and refined product demand, such as Davis and Kilian (2011), Kilian and Murphy (2014), and Coglianesi et al. (2015), find that oil demand is more elastic in the short run, with Brown (see Appendix, Section D) using a value of -0.175 in a range of -0.10 to -0.25 to represent the newer literature.

Conclusions

Regarding the six issues identified that may help put the new lower estimates of the oil security premiums into perspective, several points seem clear. Reduced oil imports likely do reduce the overall vulnerability of the US economy to world supply disruptions. Similarly, reductions in the short-run demand elasticities of oil and refined products also reduce US vulnerability. At the same time, our confidence in the new elasticity estimates, especially the elasticity of GDP with respect to oil prices, is lower than we would like it to be. The fact that the historically large oil supply disruptions are concentrated in an earlier period not well covered by the most current estimates reduces our confidence in them. Also, the evidence is still quite limited on the potential asymmetries involved—namely, the notion that increased prices may have a bigger negative effect on economic activity than the positive effect from decreased prices. For both issues, additional research is needed.

Certainly, the observed reductions in the oil-to-GDP ratio over the past 40 years are quite real, but what do these reductions imply about the sensitivity of the US economy to oil

price shocks originating from supply disruptions? While the evidence is limited, we have good reason to believe that the reduction in the oil-to-GDP ratio has contributed to the overall reduction in the GDP sensitivity. Finally, we consider the lack of modeling to address foreign oil demand shocks. Here we remain relatively confident that, short of a Fukushima-like event, changes in foreign oil demand are unlikely to be sudden. Thus new modeling in this area is not a high priority.

9. A Policy Perspective on the Oil Premium

Ultimately, the purpose of estimating the costs of US dependence on oil consumption is to provide guidance for US energy policy. The various approaches to quantifying the oil security premiums and the differing assumptions made about the elasticities can lead to substantially different estimates of the costs of US dependence on oil. Some of the estimates are consistent with relatively little intervention in US oil markets, whereas others would support considerably more intervention.

With elasticities from the older literature, the oil security premium for US consumption of imported oil averages \$6.92 per barrel over the 2015–40 time horizon. In contrast, the estimates for DSGE, SVAR, NEMS, and Benchmark-N range from \$0.39 to \$1.64 per barrel.

For US consumption of domestically produced oil, the oil security premium averages \$5.36 per barrel over the 2015–40 time horizon for Benchmark-O, with a range of \$0.17 to \$0.58 per barrel for DSGE, SVAR, NEMS, and Benchmark-N. The oil security premium for the substitution of imported for domestic oil averages \$1.57 per barrel over the 2015–40 time horizon for Benchmark-O, with a range of \$0.10 to \$0.39 per barrel for DSGE, SVAR, NEMS, and Benchmark-N.

One way of gaining perspective on the size of these premiums is to compare them with the environmental costs of US oil use. Brown

and Huntington (2015) combine estimates from a number of sources (such as Hall 1990, 2004; Fankhauser 1994; NRC 2009; Johnson and Hope 2012; US Interagency Working Group 2013; and Parry et al. 2014) to provide illustrative estimates of the environmental costs of US oil use. As shown in Table 4, replicating work by Brown (see Appendix, Section D), the resulting estimates include the

social costs of local pollution and the CO₂ emissions that result from US oil consumption.¹⁴ Estimates of oil premiums based on the older elasticities put the costs of US reliance on imported oil at roughly half the environmental costs of US oil use. In contrast, the narrower oil security estimates based on the newer elasticities are much smaller than the environmental costs of US oil use.

TABLE 4. ENVIRONMENTAL COSTS OF US OIL USE (US\$2015 PER BARREL)

Source	Environmental Costs Other Than for CO ₂ Emissions	Costs of CO ₂ Emissions
Hall (1990, 2004)	\$20.22	\$2.61
Fankhauser (1994)	n.a.	\$4.60 \$1.49 to \$10.67
NRC (2009)	\$16.79	median \$5.23 mean \$15.68 \$0.52 to \$44.42
Johnson and Hope (2012)	n.a.	\$30.58 to \$63.03
US Interagency Working Group (2013)	n.a.	\$16.32
Parry et al. (2014)	\$12.11	\$16.46

Source: Adapted from Brown and Huntington (2015).

Note: n.a. = not applicable

¹⁴ The estimated costs associated with CO₂ emissions are highly uncertain and are likely to be significantly revised by future studies. It also should be noted that a focus on the environmental costs of US oil consumption abstracts from the possibility that the environmental effects associated with production and transportation may differ between imported and domestically produced oil.

References

- Atkins, Frank J., and S. M. Tayyebi Jazayeri. 2004. "A Literature Review of Demand Studies in World Oil Markets." Discussion paper 2004-7. Calgary, AB: University of Calgary, Department of Economics.
- Balke, Nathan S., and Stephen P. A. Brown. 2016. "Oil Supply Shocks and the U.S. Economy: An Estimated DSGE Model." Unpublished manuscript, revised draft, December. Dallas and Las Vegas: Southern Methodist University and University of Nevada.
- Balke, Nathan S., Stephen P. A. Brown, and Mine K. Yücel. 2002. "Oil Price Shocks and the US Economy: Where Does the Asymmetry Originate?" *Energy Journal* 23: 27–52.
- . 2010. "Oil Price Shocks and U.S. Economic Activity: An International Perspective." Discussion paper 10-37. Washington, DC: Resources for the Future.
- Baumeister, Christiane, and James D. Hamilton. 2015. "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks." Unpublished manuscript. Notre Dame, IN, and San Diego: University of Notre Dame and University of California.
- Beccue, Phillip, and Hillard G. Huntington. 2016. "An Updated Assessment of Oil Market Disruption Risks: Final Report." Energy Modeling Forum. Stanford, CA: Stanford University.
- Blanchard, Olivier, and J. Gali. 2010. "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?" In *International Dimensions of Monetary Policy*, edited by Jordi Gali and Mark J. Gertler, 373–428. Chicago: University of Chicago Press.
- Bohi, Douglas R., and W. David Montgomery. 1982a. "Oil Prices, Energy Security, and Import Policy." Technical report. Washington, DC: Resources for the Future.
- . 1982b. "Social Cost of Imported Oil and U.S. Import Policy." *Annual Review of Energy* 7: 37–60.
- Bohi, Douglas R., and Michael A. Toman. 1993. Energy Security: Externalities and Policies. *Energy Policy* 21 (11): 1093–1109.
- Broadman, Harry G. 1986. "The Social Cost of Imported Oil." *Energy Policy* 14 (3): 242–52.
- Brown, Stephen P.A., and Hillard G. Huntington. 2013. "Assessing the U.S. Oil Security Premium." *Energy Economics* 38: 118–27.
- . 2015. "Evaluating the Costs of U.S. Reliance on Imported Oil." *Energy Policy* 79: 9–22.
- Brown, Stephen P.A., David B. Oppedahl, and Mine K. Yücel. 1996. "Oil Prices and Aggregate Economic Activity: A Study of Eight OECD Countries." Working paper 96-13. Dallas: Federal Reserve Bank of Dallas.
- Brown, Stephen P.A., and Mine K. Yücel. 1995. "Energy Prices and State Economic Performance." Economic Review, Second Quarter. Dallas: Federal Reserve Bank of Dallas.
- . 2013. "The Shale Gas and Tight Oil Boom: U.S. States' Economic Gains and Vulnerabilities." Energy brief. New York: Council on Foreign Relations.
- Coglianesi, John, Lucas W. Davis, Lutz Kilian, and James H. Stock. 2015. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." NBER Working Paper No. 20980. Cambridge, MA: National Bureau of Economic Research.

- Cooper, John C.B. 2003. "Price Elasticity of Demand for Crude Oil: Estimates for 23 Countries." *OPEC Review* 27 (1): 1–8.
- Council on Foreign Relations. 2006. "National Security Consequences of U.S. Oil Dependency." Independent Task Force Report No. 58. Washington, DC: Council on Foreign Relations Press.
- Dahl, Carol. 2010a. "Overview of Oil and Oil Product Studies." Golden, CO: Division of Economics and Business, Colorado School of Mines.
- . 2010b. "Review and Critique of Elasticities Used in the World Energy Projections Plus Model." Golden, CO: Division of Economics and Business, Colorado School of Mines.
- Davis, Lucas W., and Lutz Kilian. 2011. "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics* 26: 1187–1214.
- Davis, Steven J., and John Haltiwanger. 2001. "Sectoral Job Creation and Destruction Responses to Oil Price Changes and Other Shocks." *Journal of Monetary Economics* 48 (3): 465–512.
- DOE (U.S. Department of Energy). 2017. "Valuation of Energy Security for the United States." Report to Congress. Washington, DC: DOE.
- Edelstein, Paul, and Lutz Kilian. 2009. "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics* 56 (6): 766–79.
- . 2016. *Annual Energy Outlook 2016 (AEO2016)*. Washington, DC: US Department of Energy.
- EMF (Energy Modeling Forum). 1982. "World Oil." EMF Working Group 6 Final Report. Stanford, CA: Stanford University.
- Fankhauser, Samuel. 1994. "The Social Costs of Greenhouse Emissions: An Expected Value Approach." *Energy Journal* 15 (2): 157–84.
- Greenwood Jeremy, Zvi Hercowitz, and Gregory W. Huffman (1988) "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review* 78 (3): 402–17.
- Hall, Darwin C. 1990. "Preliminary Estimates of Cumulative Private and External Costs of Energy." *Contemporary Policy Issues* 8 (3): 283–307.
- . 2004. "External Costs of Energy." In *Encyclopedia of Energy*, vol. 2, edited by Cutler J. Cleveland, 651–67. Amsterdam: Elsevier–Academic Press.
- Hamilton, James D. 1996. "This Is What Happened to the Oil Price Macroeconomy Relationship." *Journal of Monetary Economics* 38 (2): 215–20.
- . 2003. "What Is an Oil Shock?" *Journal of Econometrics* 113 (2): 363–98.
- . 2009. "Causes and Consequences of the Oil Shock of 2007–08." *Brookings Papers on Economic Activity*, 215–61.
- . 2011. "Nonlinearities and the Macroeconomic Effects of Oil Prices." *Macroeconomic Dynamics* 15 (S3): S364–78.
- Herrera, Ana María. 2016. "The Role of Oil Supply Shocks on U.S. Economic Activity: What Have We Learned?" Unpublished manuscript, revised draft, December. Lexington: University of Kentucky.
- Herrera, Ana María, and Mohamad B. Karaki. 2015. "The Effects of Oil Price Shocks on Job Reallocation." *Journal of Economic Dynamics and Control* 61: 95–113.
- Herrera, Ana María, and Elena Pesavento. 2009. "Oil Price Shocks, Systematic Monetary Policy and the Great Moderation." *Macroeconomic Dynamics* 13 (1): 107–37.

- Hickman, Bert, Hillard Huntington, and James Sweeney, eds. 1987. *Macroeconomic Impacts of Energy Shocks*. Amsterdam: North-Holland.
- Huntington, Hillard G. 2005. “The Economic Consequences of Higher Crude Oil Prices.” Energy Modeling Forum Special Report 9. Stanford, CA: Stanford University.
- . 2016. “Oil Price Consequences of Persian Gulf Production Cuts.” Unpublished manuscript. Energy Modeling Forum. Stanford, CA: Stanford University.
- Interagency Working Group on Social Cost of Carbon. 2013. Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866, United States Government (May).
- Johnson, Laurie T., and Chris Hope. 2012. “The Social Cost of Carbon in U.S. Regulatory Impact Analyses: An Introduction and Critique.” *Journal of Environmental Studies and Sciences* 2 (3): 205–21.
- Jones, Donald W., Paul N. Leiby and Inja K. Paik. 2004. “Oil Price Shocks and the Macroeconomy: What Has Been Learned Since 1996?” *Energy Journal* 25 (2): 1–32.
- Kilian, Lutz. 2009. “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market.” *American Economic Review* 99 (3): 1053–69.
- Kilian, Lutz, and Daniel P. Murphy. 2012. “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of the Oil Market VAR Models.” *Journal of the European Economic Association* 10 (5): 1166–88.
- . 2014. “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil.” *Journal of Applied Econometrics* 29 (3): 454–78.
- Kilian, Lutz, and Robert J. Vigfusson. 2011a. “Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?” *Quantitative Economics* 2 (3): 419–53.
- . 2011b. “Nonlinearities in the Oil Price–Output Relationship.” *Macroeconomic Dynamics* 15 (S3): S337–63.
- Krichene, Noureddine. 2002. “World Crude Oil and Natural Gas: A Demand and Supply Model.” *Energy Economics* 24 (6): 557–76.
- Landsberg, Hans H., Kenneth J. Arrow, Francis M Bator, Kenneth W. Dam, Robert W. Fri, Edward R. Fried, Richard L. Garwin, S. William Gouse, William W. Hogan, Harry Perry, George W. Rathjens, Larry E. Ruff, John C. Sawhill, Thomas C. Schelling, Robert Stobaugh, Theodore B. Taylor, Grant P. Thompson, James L. Whittenberger, and M. Gordon Wolman. 1979. *Energy: The Next Twenty Years*. Cambridge, MA: Ballinger.
- Leiby, Paul N. 2008. “Estimating the Energy Security Benefits of Reduced U.S. Oil Imports.” Report ORNL/TM-2007/028. Oak Ridge, TN: Oak Ridge National Laboratory.
- Mork, Knut Anton. 1989. “Oil and the Macroeconomy, When Prices Go Up and Down: An Extension of Hamilton’s Results.” *Journal of Political Economy* 97 (3): 740–44.
- NRC (National Research Council). 2009. *Hidden Costs of Energy: Unpriced Consequences of Energy Production and Use*. Washington, DC: National Academies of Science.
- Parry, Ian W. H., and Joel Darmstadter. 2003. “The Cost of U.S. Oil Dependency.”

Discussion paper 03-59. Washington, DC:
Resources for the Future.

Parry, Ian, Dirk Heine, Eliza Lis, and Shanjun
Li. 2014. *Getting Energy Prices Right:
From Principle to Practice*. Washington,
DC: International Monetary Fund.

Smith, James L. 2009. “World Oil: Market or
Mayhem.” *Journal of Economic
Perspectives* 23 (3): 145–64.

Toman, Michael A. 2003. “The Economics of
Energy Security: Theory, Evidence,
Policy.” *Handbook of Natural Resource
and Energy Economics* 3: 1167–1218.