APPLIED SOLUTIONS TO THE CURSE OF DIMENSIONALITY IN TIME SERIES ECONOMETRICS

A Dissertation

by

ALEXANDER NEAL HLAVINKA

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	James W. Mjelde
Committee Members,	David A. Bessler
	Senarath Dharmasena
	Paul S. Busch
Head of Department,	Mark L. Waller

December 2018

Major Subject: Agricultural Economics

ABSTRACT

The curse of dimensionality is frequently encountered in applied time series econometrics when incorporating information in large datasets. Here, three applications are presented that faced challenges in dimensionality and were resolved differently: by either the a priori placement of restrictions on the parameter space or using data-driven techniques. One such data-driven method is the extraction of latent factors to reduce the information contained in many variables into a smaller number of series. A second data-driven method is statistical inference on conditional correlations to infer causality.

A parsimonious model of adoption is applied to study a large dataset of adoptions of the ductless heat pump (DHP), an energy-efficient technology. This research aims to increase the understanding of DHP adoption in the Pacific Northwest of the US by quantifying the effect of utility-provided rebates and Northwest Energy Efficiency Alliance (NEEA) expenditures on the number of installations and providing forecasts of DHP installations through 2018 given various rebate and NEEA expenditure levels. NEEA desires to increase installations of DHPs by providing funding for marketing and training; however, forecasted installations through 2018 do not meet their goals. Adoptions of DHPs are elastic with respect to the net cost of installation, and a reduction of federal tax rebates in 2011 decreases the probability of adoption.

The second objective is to investigate the dynamic effects of shocks in oil supply, aggregate demand, and oil demand on oil prices, the upstream, midstream, and downstream sectors of the petroleum industry, and the broader US economy. Because

neither the petroleum industry nor the economy can be described in a small number of series, the analysis is performed in a data-rich environment, applying a time-varying parameter extension to the factor-augmented vector autoregression. Results suggest the effects of shocks in oil supply, aggregate demand, and oil demand have evolved, and oil supply shocks play a larger role in the dynamics of the petroleum industry during recessions.

The final objective is to investigate the appropriateness of the PC Algorithm as a subset vector autoregression (VAR) methodology in determining both the contemporaneous and lag structure of the data-generating process. Subset VARs might improve forecasts and/or allow for the inclusion of more time series through a reduction in parameterization. Monte Carlo experiments show the PC Algorithm is effective at discovering the lag and contemporaneous structure of a VAR. Additional observations increase the algorithm's efficacy. When researchers do not have a priori knowledge of the true number of lags in the data-generating process, overfitting provides a smaller penalty than underfitting.

DEDICATION

To my wife, Elizabeth, my parents, Michael and Laura, and my sister, Victoria

ACKNOWLEDGEMENTS

I wish to thank Dr. Mjelde, my research advisor, for the patient guidance he has provided over last several years. What began as a research grant became a long-lasting mentoring relationship. I also wish to thank the other faculty members who served on my committee, Dr. Bessler, Dr. Dharmasena, and Dr. Busch, for their time and support while I completed this research.

I also thank my wife, Elizabeth, for her patience, endurance, and support while I completed this research.

In addition, portions of this research were conducted with the advanced computing resources and consultation provided by Texas A&M High Performance Research Computing and are tremendously appreciated.

I also wish to thank those who make their code readily available for study, especially Gary Koop and his MATLAB implementation of the time-varying parameter factor augmented vector autoregression and also the Tetrad Project team for their Java implementation of Tetrad and the PC Algorithm.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committed consisting of Dr. James W. Mjelde (advisor), Dr. David A. Bessler, and Dr. Senarath Dharmasena of the Department of Agricultural Economics and Dr. Paul Busch of the Department of Marketing.

The data used in Chapter II were provided by the Northwest Energy Efficiency Alliance (NEEA). Christine Holland was also a contributor to the research in Chapter II.

Funding Sources

Funding for the study in Chapter II was provided by NEEA, and graduate study was supported by scholarships from Texas A&M University.

NOMENCLATURE

AIC	Akaike Information Criterion
CPI	Consumer Price Index
DAG	Directed acyclic graph
DGP	Data-generating process
DHP	Ductless heat pump
DIC	Deviance Information Criterion
FAVAR	Factor-augmented vector autoregression
FEVD	Forecast error variance decomposition
FPE	Final Prediction Error
HPD	Highest posterior density
HVAC	Heating, ventilation, and air conditioning
IP	Industrial Production Index
IRF	Impulse response function
MA	Moving average
MCMC	Markov Chain Monte Carlo
MSE	Mean squared error
NEEA	Northwest Energy Efficiency Alliance
NWDHPP	Northwest Ductless Heat Pump Project
OLS	Ordinary least squares
PPI	Producer Price Index

SHD	Structural Hamming Distance
SHDR	Structural Hamming Distance per Relationship
SHDCR	Structural Hamming Distance per Contemporaneous Relationship
SHDLR	Structural Hamming Distance per Lag Relationship
SHDTR	Structural Hamming Distance per Total Number of Relationships
SIC	Schwarz Information Criterion
SIRF	Structural impulse response function
SVAR	Structural vector autoregression
TTS	Time to stabilization
TVP	Time-varying parameters
TVP-FAVAR	Time-varying parameter factor-augmented vector autoregression
VAR	Vector autoregression
2V1L	Two-variable, one-lag data-generating process
2V2L	Two-variable, two-lag data-generating process
3V1L	Three-variable, one-lag data-generating process

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

LIST OF TABLES

Page

CHAPTER I

INTRODUCTION

Bellman (1961) introduced the notion of the "curse of dimensionality." Referring originally to dynamic programming, he was describing a situation in which high-dimensional problems require substantial computing resources (Taylor 1993). Since then, computing capabilities have expanded tremendously, and the curse of dimensionality perhaps deals less with a lack of resources than it does with the shortcoming of methodologies. The curse of dimensionality may now be a viewed as a "curse of big data." There are more data available than ever, and empiricists wish to make use of a wide variety of information to address problems. Statistical analyses, however, become difficult as the number of variables increases. The problem is no longer, "Can it be feasibly done?" but rather, "How can it be feasibly done?" This is true, too, of modern time series and econometric analyses. This dissertation consists of three essays featuring methodologies dealing with large datasets.

One method of dealing with the curse of dimensionality is to impose a priori a specific functional form and restrict many potential parameters to be zero. In the first study, a large dataset on consumers' adoption of an energy-efficient technology, the ductless heat pump (DHP), is used to estimate a model of consumer adoption. This research aims to increase the understanding of DHP adoption in the Pacific Northwest of the US by quantifying the effect of utility-provided rebates and Northwest Energy Efficiency Alliance (NEEA) expenditures on the number of installations and providing

forecasts of DHP installations through 2018 given various rebate and NEEA expenditure levels. Specifically, to what extent are adoptions influenced by rebates offered by utility providers? Does funding for installer training and advertising increase adoptions? By understanding the adoption of DHPs and how adoption is affected by marketing and rebate programs, better policies can be implemented to increase installations of DHPs and reduce energy consumption (Cooney, Pater, and Meadows 2008). The adoption of innovation theory (Rogers 1962) provides the theoretical backdrop to Bass' (1969) model of durable adoption. In the proposed study, extensions of Bass' (1969) simple model by Jain and Rao (1990) and Fernandez (1999) are used to model DHP adoption in the Pacific Northwest of the US. This model reduces the need for additional parameters by imposing the sigmoid shape on the adoption profile. Results are directed to policymakers and the utility providers of the Pacific Northwest concerned with finding cost-effective methods to promote energy efficiency and reduce future load growth.

Researchers continue to develop methods for incorporating high-dimensional datasets into time series models. In Chapters III and IV, studies address two methods for dealing with issues of dimensionality in vector autoregressions (VARs).

The objective of the study in Chapter III is to investigate the dynamic effects of shocks in oil supply, aggregate demand, and oil demand on oil prices, the upstream, midstream, and downstream sectors of the petroleum industry, and the broader US economy. Field production of crude oil is an important measure of upstream activity. For the midstream, storage is explored. Refinery capacity utilization and refiners' sales of gasoline to retail outlets are studied from the downstream sector. For the broader

economy, the Industrial Production Index and Producer Price Index are studied. To incorporate a high-dimensional set of informational series, the Bernanke, Boivin, and Eliasz (2005) factor-augmented vector autoregression (FAVAR) is employed. The FAVAR is an extension of the VAR model that incorporates the information from many time series in the traditional VAR framework. This data-driven method of dimension reduction extracts a set of latent factors from the series. This study applies Mumtaz, Zabczyk, and Ellis' (2011) time-varying parameter FAVAR (TVP-FAVAR) extension to investigate how the dynamics have evolved over time. The model is estimated using Bayesian methods. Identification is performed by sign restrictions to analyze the effects of shocks in oil supply, aggregate demand, and oil demand. The findings are beneficial for decision makers both inside the petroleum industry and concerned with the general economy to understand the magnitude and duration of shocks in oil supply and demand as well as aggregate demand.

Another method of VAR modeling to address the curse of dimensionality is the subset VAR. The subset VAR reduces the dimensionality of the VAR's estimation by eliminating regressors. The objective of the final study is to investigate the appropriateness of the PC Algorithm (Spirtes, Glymour, and Scheines 2000) as a subset VAR methodology in determining both the contemporaneous and lag structure of the data-generating process (DGP). This, too, is a data-driven technique, but instead of extracting common features of the series, the data are used to test statistically conditional correlations to infer causality. The final study evaluates, using Monte Carlo techniques, the Akleman, Bessler, and Burton (1999) directed acyclic graph-adaptation of Hsiao's

(1979) search methodology for the placement of VAR coefficient restrictions. The PC Algorithm is evaluated for its ability to detect VAR structures considering different residual volatility levels, different numbers of observations in the dataset, different numbers of lags included for the algorithm's search, and different significance levels for the hypothesis tests of statistical independence. The results are directly applicable to researchers using the VAR method 1) to reduce the parameterization; 2) to identify the "correct" model for the specific DGP; or 3) to improve forecast performance using subset VARs (Akleman, Bessler, and Burton 1999; Bruggeman 2004).

CHAPTER II

FORECASTING THE ADOPTION OF RESIDENTIAL DUCTLESS HEAT PUMPS[1]

Some of today's energy-efficient appliances consume less than half the energy consumed by their predecessors (National Resources Defense Council 2010). Ductless heat pumps (DHPs), developed in the 1970s in Japan, offer increased efficiency in heating and cooling homes (Swift and Meyer 2010). A ductless system is comprised of an outdoor and an indoor unit to distribute air for both cooling and heating. A line running between the two units requires only a three-inch hole and eliminates the need for expansive ductwork (Northwest Ductless Heat Pump Project 2014). Unlike other types of heat pumps, including geothermal heat pumps, DHPs are relatively easy and inexpensive to install (Sutherland 2012). Benefits of DHPs to homeowners include increased comfort, a reduction in electricity consumption, the ability to heat and cool with a single appliance, relatively low-cost installation, and potential financial incentives for installation, including federal and state income tax credits and utility-provided rebates (Northwest Ductless Heat Pump Project 2014). Swift and Meyer (2010) and Bugbee and Swift (2013) note almost all residential heating, ventilation, and air conditioning (HVAC) systems in Asia and the majority of those in Europe are ductless, but DHPs represent less than one percent of HVAC systems in the US. Awareness of

---

[1] Reprinted with permission from "Forecasting the Adoption of Residential Ductless Heat Pumps" by Alexander N. Hlavinka, James W. Mjelde, Senarath Dharmasena, and Christine Holland. *Energy Economics* 54: 60-67, Copyright © 2015 by Elsevier Ltd. DOI: https://doi.org/10.1016/j.eneco.2015.11.020

DHP technology in the US has increased since 2006 when redesigned, more efficient, and advanced-controlled ductless technologies were made available (Storm et al. 2012).

The Northwest Ductless Heat Pump Project (NWDHPP), a collaboration between the Northwest Energy Efficiency Alliance (NEEA) and its utility and energy partners, was established in 2008 to accelerate DHP installations in electricity-heated homes in the Pacific Northwest of the US (NEEA 2013; NWDHPP 2012). NEEA (2013) estimates the 13,000 DHPs installed in the Northwest through 2011 saved 40.5 million kilowatt hours of electricity per year. These savings represent nine percent of the potential regional savings estimated by Cooney, Pater, and Meadows (2008). This research aims to increase the understanding of DHP adoption in the Pacific Northwest of the US by quantifying the effect of utility-provided rebates and NEEA expenditures on the number of installations and providing forecasts of DHP installations through 2018 given various rebate and NEEA expenditure levels. Because DHPs were introduced relatively recently into the regional market for residential HVAC systems, the adoption of innovation theory is applied.

Literature Review

Energy-Efficient Technologies and Consumers

Vast amounts of public and private capital have been used to advocate energy efficiency. Despite the many energy-efficient appliances available to today's consumers, studies show that consumers do not always adopt the economically feasible set of durables (Gates 1983; Howarth and Andersson 1993; Howarth and Sanstad 1995). Energy-efficient innovations' up-front costs, including the costs of information and of the good

itself, are relatively high compared to those of conventional technologies (Howarth and Sanstad 1995). Gates (1983) argues that even though the returns for conservation investments such as setting back the thermostat, weather stripping, caulking, and adding insulation may exceed the returns of financial instruments, consumers still do not invest in energy-efficiency. Howarth and Sanstad (1995) and Gates (1983) both believe that energy-efficient technologies' low rates of adoption are the result of consumers' relatively high discount rates. A high discount rate may be evidence of the high cost of information or perceived risk of investment in these technologies.

Income and lack of knowledge of household energy consumption are barriers to compact fluorescent light adoption (Mills and Schleich 2010). Panzone (2013) estimates the demand for refrigerators, washing machines, televisions and light bulbs using Deaton and Muellbaurer's (1980) Almost Ideal Demand System and draws four major conclusions. First, the influence of own-price on purchases of energy-efficient appliances depends of the possibility of behavioral adjustments associated with the good. In addition, current energy prices may not be driving adoption as much as theory suggests. Next, energy-efficient appliances are perceived as necessities. Finally, consumers may value the good's efficiency less than they value other attributes. His final conclusion is supported by Mills and Rosenfeld (1996), who find non-energy benefits often motivate consumers to adopt an energy-efficient technology. Two non-energy benefits of compact fluorescent lights and light-emitting diodes, for example, are their reduced heat generation and longer lives. As another example, insulated windowpanes offer more comfort than non-insulated windows. Further, consumers may

adopt energy-saving technologies for a "warm-glow" feeling of promoting energy and environmental conservation. Examining the effect of residential HVAC systems characteristics on the homeowner's adoption decision, Michelsen and Madlener (2012) find the system's attributes are more relevant to owners of newly-built homes than to owners of existing homes.

Lund (2006), employing a methodology related to our own, estimates the diffusion of new energy technologies. Using international data on twenty technologies, he finds the time required for new energy technologies to reach at least 50 percent of their market potential ranges between 10 and 70 years. End-use and energy-saving consumer goods, on average, take less than 25 years to reach this level. Shorter times are indicative of a good's relatively high impact on energy production or consumption. Traditional, non-ductless heat pumps, for example, are estimated to take between 35 and 65 years to attain 50 percent of their market potential in three European countries. Lund (2006) also finds countries with subsidized energy-efficiency programs, other things equal, had higher penetration rates for the technologies relative to other countries.

The importance of financial incentives to potential adopters of energy-efficient goods has been the subject of several other studies. Using a choice experiment, Aalbers et al. (2009) find that a subsidy may entice firm managers to adopt a technology even if the subsidy is too small to make the technology profitable. It may be that "…the presence of a subsidy invokes a positive connotation…[that] may carry enough weight in an agent's decision making to tip the balance in favor of the subsidized technology" (Aalbers et al. 2009, p. 439). Wasi and Carson (2013) examine the role of rebates in

shifting the percentage of electric water heaters to climate-friendly units in New South Wales, Australia. The rebate program increases the number of climate-friendly heaters in homes without access to natural gas. Murray and Mills (2011) conclude that the effect of rebates offered for certain Energy Star technologies as part of the American Recovery and Reinvestment Act of 2009 is indeterminate and the widely-available rebates were quickly exhausted. Instead, rebates that target the marginal consumer would have had more of an impact on purchases of energy-efficient goods. Rebate policies, however, "… may also encourage large-scale purchasing of energy-efficient appliances, which may finally result in an increase in electricity consumption (rebound effect)" (Galarraga, Abadie, and Ansuategi 2013, p. S98).

Many studies have examined consumers' investments in energy-efficient technologies, but few have studied the market-level adoption of these goods. Several studies have examined DHPs, but the majority of studies has been engineering-oriented; see, for example, Fransisco et al. (2004), Şahin, Kılıç, and Kılıç (2011), and Stecher and Allison (2012). Cooney, Pater, and Meadows (2008) note that regional electricity providers are seeking to meet a greater portion of load growth through energy efficiency. They suggest DHPs have achievable savings of upwards of 438 million kilowatt hours per year for the Pacific Northwest of the US, but at the time of their publication, there were few installers with knowledge of the product. Based on 144 installed DHPs in a pilot study in Connecticut and Massachusetts, annual household energy savings of approximately $400 are reported (Cooney, Pater, and Meadows 2008). Storm et al. (2012) discuss a pilot program developed by NEEA as a precursor to the NWDHPP.

This pilot program successfully increased consumers' awareness and interest in DHPs. Through supplier training and distribution networks, NEEA also improved the supply-side of the market. Overall, Storm et al. (2012, p. 2-304) conclude DHPs are "…an important and transformational technology."

Literature on Adoption Models

Because DHPs were introduced relatively recently into the Northwest regional market for residential heating, the adoption of innovation theory is applied. Models of innovation adoption have a common root in the sociological work of Rogers (1962), who identifies four elements of the diffusion process: innovation, communication channels, time, and the social system. Theory and empirical work show that that the adoption profile, a plot of cumulative adopters over time, is sigmoid-shaped. Diffusion depends on social and economic factors that vary over time and among different groups of potential adopters. Rogers' (1962) original model imposes a Normal distribution upon the time of adoption across consumers described as either innovators, early adopters, early majority, late majority, or laggards.

Models of innovation adoption have extended Rogers' (1962) pioneering work. One of the first mathematical extensions of the adoption model is Bass (1969). In this model, the instantaneous rate of adoption *f* at time *t* is given by the differential equation

$$f(t) = \left[\alpha + \beta F(t)\right]\left[1 - F(t)\right] \tag{2.1}$$

where α is the coefficient of innovation, β is the coefficient of imitation, and F is the proportion of all adopters who have adopted by time t. The coefficient of innovation gives the probability of purchase when t is zero and captures the innovativeness of

potential adopters. The coefficient of imitation "…reflects the pressure operating on imitators as the number of previous buyers increases" (Bass 1969, p. 216).

Empirical work using the Bass (1969) model finds that adoption is explained well without any other variables. The model has been used effectively in the retail, industrial, and agricultural sectors (Bass, Krishnan, and Jain 1994). Extensions of Bass' basic model allow for the inclusion of explanatory variables intended to improve managers' marketing decisions related to influencing either the rate of adoption or market potential. Jain and Rao (1990) and Fernandez (1999) allow the adoption curve to be shifted by a vector of exogenous variables and also estimate demand elasticities.

Why Subsidize the Adoption of DHPs?

In 2011, 92 utility companies in the Pacific Northwest of the US offered some form of a subsidy for the installation of DHPs. Economic theory suggests that incentives can increase the adoption of energy-efficient innovations. There are at least two commonly cited reasons to subsidize the use of electricity-conserving innovations (Howarth and Andersson 1993; Jaffe and Stavins 1994; Levine et al. 1995). The first is the existence of market failures. Failures in energy markets may include environmental externalities, average-cost electricity pricing, national security issues, liquidity constraints, and information problems (Gillingham, Newell, and Palmer 2009). Even if the market is functioning correctly, there may still be reasons to incentivize. For utility providers faced with increasing demand for electricity, a subsidy promoting electricity conservation may help defer costly expansion projects (Loughran and Kulick, 2004; Tietenberg and Lewis 2012). Adding new generating facilities or bringing a facility with

a high marginal cost of operation online is a costly endeavor that may lead to higher generating costs. Further, new plants are designed with larger generating capacities than necessary to satisfy the current demand. As such, there is unused capacity, resulting in increased costs. By subsidizing energy-efficiency, utilities delay the need for additional capacity.

Theoretical Model

Extensions of Bass (1969) by Jain and Rao (1990) and Fernandez (1999) are the basis for the adoption model. The exposition of the theoretical model follows from Jain and Rao (1990, pp. 164-165). Solving the differential equation in equation (2.1) gives

$$F(t) = \frac{1 - \exp\left[-(\alpha + \beta)t\right]}{1 + (\beta/\alpha)\exp\left[-(\alpha + \beta)t\right]}. \tag{2.2}$$

The proportion of the market potential that adopts in the interval $(t-1,\ t)$ is

$$F(t) - F(t-1). \tag{2.3}$$

Assume an individual's probability of adoption is π. The expected proportion of the total market that adopts between $t-1$ and t is

$$\pi\left(F(t) - F(t-1)\right). \tag{2.4}$$

The conditional probability of adoption in the same interval given the consumer has not yet adopted is

$$\frac{\pi\left(F(t) - F(t-1)\right)}{1 - \pi F(t-1)}. \tag{2.5}$$

If A_t represents the total number of adoptions by time t and M is the market potential representing the upper limit on the potential number of adoptions, sales, S_t, in the interval $(t-1,\ t)$ are

$$S_t = (M - A_{t-1})\frac{\pi\left(F(t) - F(t-1)\right)}{1 - \pi F(t-1)}. \tag{2.6}$$

Let the probability of adoption vary over time according to a logistic function of a vector of covariates

$$\pi(x_t) = \frac{\exp\left[\gamma \cdot \ln x_t\right]}{1 + \exp\left[\gamma \cdot \ln x_t\right]} \in (0,1), \tag{2.7}$$

where x_t is a k-element vector of variables and γ is a k-element vector of coefficients interpreted as the elasticity of adoption with respect to the variable. While Jain and Rao (1990) use only the price of the good, Fernandez (1999) includes a vector of exogenous variables in addition to the good's price. Both studies interpret the coefficients on the natural logarithm of each variable as the elasticity of demand with respect to the given variable.

Data

Data on DHP sales, provided by NEEA, consist of 15,606 observations of single-family households' purchases of DHPs in Oregon, Washington, Idaho, and Montana. The data span 58 months from January 1, 2009 through August 30, 2013. Available data include the number of units installed, the date of installation, costs of equipment and installation, NEEA program expenditures, and a rebate schedule. All nominal dollar amounts have

been adjusted to February 2014 dollars using the consumer price index (Federal Reserve Bank of St. Louis 2018) for all items in the US.

Each observation is one participant in the NWDHPP. Participating households must meet certain eligibility requirements. First, the house's primary heating appliance must be powered by electricity. New construction, homes with fossil fuel forced-air or hydronic heating systems, and houses with existing DHP installations are also ineligible to participate in the program. For this analysis, only outdoor DHP units are considered because more than one air-handling unit can be installed per outdoor unit, but all outdoor units must have at least one air-handling unit. The number of outdoor units "better" reflects the heating and cooling needs of the house. In addition, only unique installations are considered. If a household installed multiple units in a given installation, they are considered a single adoption. From this point forward, an installation will refer to the adoption of a DHP system – regardless of the number of indoor or outdoor units – by a single-family household heated by electricity in Oregon, Washington, Idaho, or Montana.

The installation data are aggregated to form a monthly dataset. Summary statistics are presented by month in Table 2.1. Installations generally increase as the year's end approaches. Annual installations increase until 2011 when the federal tax credit for DHP installations, originally set to expire in 2010, was extended but decreased from a maximum of $1,500 to $300 (Evergreen Economics 2012). Consumers may have elected to install a DHP before the incentive was reduced. Historical monthly DHP installations are shown in the unshaded portion of Figure 2.1. Installations of DHPs

exhibit seasonality. While the peak number of installations occurs in December 2010, there are smaller peaks in December of the other years, too. The increase in sales in December may be a result of lower temperatures or homeowners taking advantage of tax incentives before the year's end.

The total cost for a given installation is the sum of the costs of equipment, labor, electrical, tax, and other costs. Real installation costs for DHPs are fairly constant throughout a given year but are increasing over the entire period. The average real per-unit DHP installation cost increased approximately 3.7 percent annually or roughly 20 percent from 2009 through August 2013.

From its beginning through October 2013, the NWDHPP provided over $9 million for advertising, education, and training for installers (NEEA 2013). DHP advertising campaigns include the NWDHPP website, radio, TV, and print ads and public service announcements, flyers for distribution by HVAC installers, and door hangers. Training for HVAC installers includes webinars and Master Installer certifications. Data on NEEA's monthly expenditures are available beginning in January 2010; they are estimated for 2009 using the total expenditures for 2009 and each month's average proportion of annual expenditures from the subsequent years. NEEA's monthly expenditures reflect the pattern of sales of DHPs; both have yearly peaks in December. Annual expenditures have generally been increasing.

Empirical Models

Installation Model

The probability of installation by a household during month *t*, shown in equation (2.7), is modeled as a function of several variables. The first is the net per-unit cost of installation, that is, the per-unit total cost for a given household's installation of a DHP less all federal and state tax credits and utility-provided rebates. The median net per-unit installation cost for all installing households in a given month is used. Contemporaneous NEEA expenditures on marketing also affect the probability of adoption. Indicator variables for month control for seasonality. Finally, an indicator variable captures the change in the federal tax code. The installation model is

$$S_t = (M - A_{t-1}) \frac{\hat{\pi}_t \left(F(t) - F(t-1) \right)}{1 - \hat{\pi}_t F(t-1)} + \hat{\varepsilon}_t, \tag{2.8}$$

where $\hat{\pi}_t$ is

$$\hat{\pi}_t = \frac{\exp\left[\hat{c} + \sum_{i=1}^{2} \hat{\varphi}_i \ln(x_{it}) + \sum_{w=1}^{11} \hat{\delta}_w g_{wt} + \hat{\gamma} I_{tax}(t) \right]}{1 + \exp\left[\hat{c} + \sum_{i=1}^{2} \hat{\varphi}_i \ln(x_{it}) + \sum_{w=1}^{11} \hat{\delta}_w g_{wt} + \hat{\gamma} I_{tax}(t) \right]} \tag{2.9}$$

with variables:

- S_t the number of installations, specified in single-family households heated by electricity, in month *t*;
- *M* the total market potential of single-family households heated by electricity;
- A_t the cumulative number of installations in month *t*;
- *F* the differential equation as defined in equation (2.1);

- x_{1t} the median net cost per unit (in dollars per unit) to install a DHP in month *t*;
- x_{2t} NEEA's expenditures on the NWDHPP (in dollars) in time *t*;
- g_{wt} indicator variable for month *w* in time *t* such that January is the base;
- $I_{tax}(t)$ an indicator function that equals one for all months prior to 2011, the time of the change in the tax code, and 0 otherwise;
- $\hat{\varepsilon}_t$ the error term; and
- $\hat{c}, \hat{\varphi}_i, \hat{\delta}_w$, and $\hat{\gamma}$ are estimated parameters.

Unfortunately, the dataset does not include household income. Instead, the model was estimated using the median Zillow-estimated value of all houses that installed a DHP in a given month as a proxy for income and a regional, consumption-weighted average of state-level retail electricity prices. The model presented here performed better than models with the income proxy and electricity price. Given the previously-discussed findings of Panzone (2013), current electricity price may not be a driver of DHP adoption.

The choice of estimation techniques used in the adoption of innovation literature has evolved. Bass (1969) employed ordinary least squares (OLS). Both the maximum likelihood estimator and nonlinear least squares estimation techniques models provide considerably improved step-ahead forecasts compared to the OLS models (Schmittlein and Mahajan 1982; Srinvasan and Mason 1986). Srinivasan and Mason (1986), however, show that the maximum likelihood parameters may be misestimated because of the misspecification of the distribution function of adoption time and that find

nonlinear least squares estimated models provide a better fit – measured by mean squared error and mean absolute deviation – in most cases compared to models estimated using maximum likelihood. Equations (2.8) and (2.9), therefore, are estimated using nonlinear least squares.

Starting values for the coefficients of innovation and imitation were selected from the OLS estimation of the simple Bass (1969) model. The starting value for the market potential, 500,000 single-family households, is based on discussions with NEEA personnel; estimation of the market potential is restricted to lie between the number of total installations and the number of single-family households heated by electricity in the region. The parameter on net DHP installation cost is restricted to be negative. To guarantee an appropriate shape for the adoption profile, the coefficient of innovation must be non-negative to have a positive adoption profile. Estimations were performed using SAS 9.3 (SAS Institute 2013).

Simulation Model

A stochastic simulation model based on distributions for both the error term and estimated parameters from the installation model is developed to generate probabilistic forecasts. The estimated parameters are assumed to be distributed by a truncated multivariate Normal distribution centered at the vector of estimated parameters estimated from the empirical model in equations (2.8) and (2.9). The covariance matrix of the estimated parameters gives the distribution's covariance matrix. Truncation of the distribution is necessary for three reasons. First, the market potential cannot exceed the number of households in the target region, nor can it be less than the total number of

installations at the beginning of the forecast period. Second, the parameter on net installation cost must be negative. Lastly, the coefficient of innovation must be greater than or equal to zero. These truncations correspond to the restrictions imposed at estimation.

Monthly installations are forecasted using the installation model of equations (2.8) and (2.9). In the stochastic simulation, total per-unit installation cost before any rebates or tax credits are drawn from a Normal distribution with mean and variance equal to their sample counterparts across all observations. Any applicable rebates and tax credits are then subtracted from the total installation cost to find the net installation cost. Each month's installations are added to the previous month's cumulative number of installations $A_t = S_t + A_{t-1}$. Let z_t be the rebate amount. Rebates awarded in time t are equal to the product $S_t z_t$. Total yearly rebates are the sum of the rebates awarded in each month for a particular year. Annual rebates awarded can be no more than the predetermined annual rebate budget. Once the rebate budget is reached, the rebate for any further installations in that year is zero. All simulations are performed using R version 3.1.1 (R Core Team 2014).

Results

Estimation Results

A Breusch-Godfrey test (Breusch and Pagan 1980) for serial correlation shows significant first-order autocorrelation, and visual inspection of the autocorrelation function of residuals reveals it slowly decays, indicating the presence of an autoregressive process for the error. The model is refit assuming an AR(1) process for

the error term. Parameter estimates and their associated p-values are displayed in Table 2.2. Throughout the discussion, the five percent level of significance is assumed.

The estimated market potential is 329,442 single-family, electricity-heated households, however, its p-value is 0.395. The coefficient of innovation is estimated to be 0.002, but it is not statistically significant; this estimate is reasonable given the slow increase in monthly sales early in the dataset. The coefficient of imitation is 0.068 and is statistically significant. Once DHPs are established in the market, consumers may be more compelled to install them, but cumulative installations will increase slowly at first. As previously discussed, this is not an uncommon phenomenon for energy-efficient technologies; see Howarth and Sanstad (1995) and Gates (1983). A meta-analysis of 213 studies using the Bass model finds the average coefficients of innovation and imitation are 0.03 and 0.38, but the coefficient of imitation varies substantially between studies (Sultan, Farley, and Lehmann 1990). In another meta-analysis, Van den Bulte and Stremersch (2004) use the natural logarithm of the coefficient of imitation divided by the coefficient of innovation to compare studies. The average in their sample is 3.42 with a standard deviation of 2.13; the estimated ratio for this analysis is 3.526.

The elasticity of installations with respect to net installation cost is approximately -1.495 and is statistically significant. NEEA expenditures are estimated to have a positive, inelastic effect on demand. The estimate is 0.316 but is not statistically significant. The monthly indicator variables reproduce the visual pattern seen in seasonal sales; coefficients on months late in the year are significant and relatively larger. A structural break appears to be associated with the tax code change as shown by

the marginally statistically significant indicator variable. Evaluating equation (2.9) when $I_{tax}(t)$ is equal to zero and then equal to one and at the mean of all other variables gives a probability of installation before 2011 that is over two times the probability of installation beginning in 2011.

Forecasted Installations through 2018

All assumptions for the simulations are given in Table 2.3. Annual proposed budgets for 2014 through 2018 for the NWDHPP rebates and NEEA's expenditures are provided by NEEA. All dollar amounts are assumed to be in real dollars. Each month's NEEA funding is assumed to follow its historical proportion of annual expenditures. Using these projections, the five-year expected number of DHP installations by month are shown along with historical installations in Figure 2.1. The forecast mirrors the intra-year seasonality shown in the historical data and shows a gradual decrease in the rate of DHP installations. The cumulative distribution function (CDF) for five-year installations is shown in Figure 2.2. The expected number of installations through 2018 is 9,998, or two-thirds the number of installations in the first five years of the program. The CDF shows that the distribution is not completely symmetric; there is more area in the right tail of the density function. This demonstrates some potential for low-probability events with many additional installations. It is also important to remember that because observations with missing data on the number of outdoor units or the models of those units were dropped, the forecasts presented are conservative.

There are four potential reasons for the decrease in installations. First, DHP installations appear to be slowing. This can be seen by comparing each month's peak

from year to year beginning in 2010 in Figure 2.1. A second reason is the aforementioned increasing per-unit cost of installation. The estimated probability of installation, as shown in equation (2.9), averages approximately five percent after the decrease in the tax credit; this is primarily the result of the elastic estimate on net installation cost. Yet another reason is that the NWDHPP promotes DHPs not only as an energy-efficient form of space conditioning but also as a supplement to existing heating to enhance comfort. Consumers may be less likely to adopt if they view DHPs as purely comfort-enhancing and not as an energy-efficient technology with the potential for cost savings. Also, the NWDHPP began during the global recession of 2009. The poor economic conditions may have been inhibitive for growth in the market for DHPs.

A sensitivity analysis using rebates from $0 to $3,500 in increments of $250 is performed to quantify the rebate amount that will maximize the number of additional installations through 2018 subject to the annual budget. The results are shown in Figure 2.3; also included is a 95-percent confidence interval. As the rebate amount increases, expected installations also increase until the rebate is $2,750. The upper bound of the confidence interval, however, is already decreasing at this amount, but the lower bound is increasing. A rebate of $2,750 would increase five-year installations approximately 18 percent relative to the currently-offered rebate of $1,500. Rebates above $2,750 exhaust the rebate budget rapidly; there are a number of consumers who are unwilling or unable to pay the entire cost of the DHP less the tax credit. The 95-percent confidence interval shows that at the current rebate level, there is less certainty in the number of

installations than at a rebate of $2,750, the rebate amount that maximizes the number of expected installations through 2018.

A second sensitivity analysis is conducted by varying NEEA's annual expenditures (Figure 2.4). Annual expenditures range from $1 million to $3.5 million with every year in the five-year forecast receiving the same funding. Unlike the sensitivity analysis on rebates, this analysis shows that as NEEA's expenditures increase, expected installations also continue to increase. Expected installations change less noticeably in response to changes in NEEA's expenditures than to changes in the rebate amount. The slow rate is the result of the inelastic effect of NEEA's expenditures on installations, while the effect of the rebates is directly "felt" by consumers because of the immediate decrease the net cost of installing a DHP. The upper bound of the 95 percent confidence interval demonstrates the potential for a substantial increase in the number of installations given additional funding.

At first glance, the original five-year installation forecast and the sensitivity analysis performed on different levels of NEEA expenditures seem to be in contradiction. The historical annual average of NEEA's expenditures is just under $2 million. Given the analysis in Figure 2.4, this level of expenditure is expected to generate approximately 11,335 additional installations, but this represents 13 percent more installations than suggested by the "status-quo" forecast. This discrepancy occurs because each year receives the same funding in the sensitivity analysis on NEEA's expenditures. In the five-year forecast, the projected budgets for NEEA, shown in Table 2.3, vary from year to year and are usually less than the historical average. This

difference is important because of the sigmoid-shaped adoption profile; more funding earlier in the adoption process will increase the probability of adoption and the total number of installations.

A final sensitivity analysis is performed by varying both NEEA's expenditures and the rebate amount (Figure 2.5). This simulation a shows the tradeoffs between NEEA expenditures and rebates. As an example of reading the figure, if the offered rebate and NEEA's expenditures are $2,500 and $3 million, then the expected number of additional installations is 12,448. The lower bound of the confidence band is 1,987 additional installations, but the upper bound is 24,011. The effect of the limited rebate budget is seen where both the average and upper limit of the confidence interval decrease after a rebate amount of approximately $2,500.

The surfaces support the results of the single-variable sensitivity analyses. First, rebates are more effective at increasing additional installations than are NEEA's expenditures. The upper bound of the confidence interval shows the potential, for a substantial number of new installations relative to the number installed through August 2013. As an example, policymakers can achieve approximately 11,000 additional installations from 2014 through 2018 by offering rebates of $2,500 per installation along with $1.1 million in marketing per year, rebates of $2,000 with $1.6 million for marketing, or even the current rebate of $1,500 with annual marketing expenditures of $2.8 million. These changes demonstrate how sensitive the forecasted installations are to the offered rebate.

Discussion and Conclusions

Installations of DHPs are net-cost elastic; changes in the rebate have a significant effect on the expected number of installations over the next five years. In addition, NEEA's expenditures have a small, positive effect on DHP installations. Historical data, the estimated model, and subsequent forecasts show that installations of DHPs have grown slowly and the overall rate of installations will decrease without some change in the exogenous variables. A rebate of $2,750 would maximize installations through 2018 subject to the proposed rebate budgets. This amount would almost double the rebate currently offered by most utilities. It is important to remember that installations refer to those adoptions of DHPs in single-family households heated by electricity that were eligible to participate in the NWDHPP. There have also been installations of DHPs in multi-family units and households that do not use electricity for heating. Neither of these situations is considered here.

Notably, the small and statistically insignificant coefficient of innovation shows that the adoption profile for DHPs grows slowly. There are several explanations for the slow rate of growth in installations of DHPs. The most likely barrier to an increased rate of installation is up-front costs. The literature on adoption of innovation and consumers' purchases of energy-efficient technologies shows that cost is a significant obstacle to overcome, especially given the 80 percent reduction in the federal tax credit in 2011. A second reason, as discussed by Golder and Tellis (2004), is that the growth rates of leisure-enhancing products, such as TVs and stereos, are higher compared to those of other innovations. Because DHPs do not fit the "leisure-enhancing" category, a slow

growth rate may be expected. NEEA has also noted that a substantial proportion of installations are occurring in households with older occupants, although Rogers (1962) argues there is inconclusive evidence on differences in age between early and late adopters. The relatively high age may be a result of the NWDHPP's promotion of DHPs as comfort-enhancing. Finally, poor economic conditions may retard adoption. The installations studied here began at the same time as the economic recession of 2009. While the market potential is large, forecasted sales are relatively small and show no growth compared to the first five years of the program. Because of their direct effect, rebates are more effective than NEEA expenditures at increasing DHP installations. If utility providers wish to make their DHP programs more successful, they must focus on the costs – direct and indirect – faced by potential adopters of DHPs. Future research should investigate how incented installations affect the installations of non-incented DHPs and investigate the adoption of other energy-efficient technologies. Is the average adoption profile for energy-efficient innovations slow-growing? Are consumers responsive to changes in prices of these goods?

Limitations exist when using the adoption of innovation theory and attempting to model early in the adoption process. Data limitations are the primary concern. It has been demonstrated that the structural validity of the model does not guarantee accurate forecasts with limited data, especially when only a small proportion of the potential market has adopted the innovation (Heeler and Hustad 1980). If indeed only a small proportion of the DHP market potential has installed, then the forecasts may not be accurate. Bass (1969) also notes that forecasts are sensitive to the coefficients of

innovation and imitation; small changes in either of these coefficients may lead to substantially different forecasts. He also notes the importance of having already observed the peak of instantaneous adoptions, that is, the point in time that maximizes equation (2.1), to forecast accurately many periods out-of-sample. There are, of course, always data limitations in the form of limited observations when forecasting the sales of a new product.

A second set of limitations pertains to the stochastic simulation and sensitivity analyses. There is uncertainty in the marketing and rebate budgets provided by NEEA. If these budgets, especially the rebate budget, are altered, the forecasts will change. In fact, some of the simulation iterations showed the possibility of exhausting the annual rebate budgets in 2015 and 2016 under the current rebate schedule. If steps are taken to prevent this, then it is possible that there may be more installations over the next five years. In addition, because NEEA's expenditures and the rebate amount do not fluctuate as dramatically in reality as they are adjusted in the sensitivity analyses, it is possible that the problems of out-of-sample forecasting may obscure the results. The relationships estimated by the installation model may hold outside of a reasonable sampling region.

CHAPTER III

AN INVESTIGATION OF OIL PRICES, THE PETROLEUM INDUSTRY, AND THE US ECONOMY

The effects of changes in oil prices are not limited to the energy industry. In modern, energy-dependent economies like that of the US, the petroleum industry's impacts are far-reaching. The oil and natural gas industry accounted for six percent of total employment, six percent of labor income, and eight percent of value added to the American economy in 2011 (American Petroleum Institute 2013).

The objective of this study is to investigate the dynamic effects of shocks in oil supply, aggregate demand, and oil demand on oil prices, the upstream, midstream, and downstream sectors of the petroleum industry, and the broader US economy. The relationships are examined over the period following the Great Moderation, a decline in business cycle volatility beginning in the early 1980s (Stock and Watson 2002b). No single measure completely describes the activity of each sector of the petroleum industry or of the entire economy, but using the information contained in many time series is not straightforward. The vector autoregression (VAR), a model commonly applied in multivariate time series analysis and forecasting, is difficult – if not impossible – to apply in situations when there are many time series. Additional series necessitate the estimation of many additional parameters. A method of dimension reduction in VAR analysis is the factor-augmented vector autoregression (FAVAR) proposed by Bernanke, Boivin, and Eliasz (2005). Extensions to the FAVAR (Mumtaz, Zabczyk, and Ellis

2011) which allow for time-varying parameters (TVP) and volatility are applied here to estimate changes among relationships over time.

The TVP-FAVAR model is applied to a large dataset comprised of series related to the petroleum industry and macroeconomic variables which capture general economic conditions. Following the literature, three proposed structural shocks, oil supply shock, aggregate demand shock, and oil demand shock, are analyzed. This study advances the literature following Baumeister and Peersman (2013), who similarly study the evolution of oil prices using Primiceri's (2005) TVP-VAR, by performing similar analyses in a data-rich setting.

The Effects of Oil Prices on the Economy

Ederington et al. (2011) provide a comprehensive review on oil prices and the factors affecting them. Oil prices and their effects on the economy are frequently studied. Hamilton (2009, p. 1) provides a succinct description of the history of oil prices: "…changes in the real price of oil have historically tended to be (1) permanent, (2) difficult to predict, and (3) governed by very different regimes at different points in time." Almost every recession between World War II and the early 1980s was preceded approximately nine months by an increase in oil prices (Hamilton 1983), and the reduction in oil prices between 2007 and 2008, at the beginning of the Great Recession, is attributed to increased global demand amid slowing supply (Hamilton 2009). While the reduction in oil prices from 2014 to 2016 is sometimes attributed to US production in shale plays, Prest (2018) finds no evidence of increased shale production causing the decline in oil prices. Instead, Baumeister and Kilian (2016) attribute over 20 percent of

the 2014 decline in oil prices to negative demand shocks, which Baumeister and Kilian (2017) later show provided a stimulus to real GDP growth. This stimulus, however, was offset by a reduction in investment in the oil industry. Lippi and Nobili (2012) find a decrease in global oil production decreases US industrial production, while Kilian (2009) shows innovations to oil supply, aggregate demand, and oil demand, have different effects on the US economy. On the other hand, oil prices are affected most by economic activity (Wang and Sun 2017). Crude oil prices are affected by changes in economic activity as measured by the Kilian Index (He, Wang, and Lai 2010).

While shocks in the price of oil may contribute to recessions, some believe fluctuations in oil prices affect the US economy less than commonly believed (Barsky and Kilian 2004). Herrera and Pesavento (2009), whose findings are supported by Baumeister and Peersman (2013) and Aastveit (2014), argue the effect on the US economy of shocks in oil prices has been substantially dampened since the Great Moderation. Kilian and Vigfusson (2017) find a one standard deviation shock in oil prices at the trough in 2008 would have had no effect on real GDP growth, while the same shock at the peak of oil prices in 1993 would have caused a persistent reduction in economic growth of three percentage points.

Employment is affected by changes in oil prices. Herrera, Karaki, and Rangaraju (2017), using the FAVAR, find decreases in net employment for oil and gas extraction and mining support within the first year following a negative shock in oil prices as jobs move to the construction, manufacturing, and services sectors. Additional rig activity

creates positive employment benefits of at least 37 jobs immediately and 224 in the long run (Agerton et al. 2014).

Oil prices are influenced by exchange rates and interest rates (Krichene 2006). A shock in oil prices leads to an increase in interest rate over several years. Similarly, Cuaresma and Breitenfellner (2008) show forecasts for oil prices are improved substantially if information on exchange rates are included. Granger causality tests for oil price and exchange rate movements, however, are inconclusive. Oil price is a good predictor of dollar-related exchange rates (Lizardo and Mollick 2010). An increase in oil prices causes the dollar to depreciate against currencies of countries with positive net oil exports. Bernanke, Gertler, and Watson (1997) believe monetary policy could be effective at countering the consequences of oil price shocks. Anzuini, Lombardi, and Pagano (2010) show expansionary shocks in US monetary policy tend to increase commodity prices. Oil prices respond sharply to shocks in the federal funds rate, money supply, the consumer price index, and an industrial production index. Expansionary monetary policy increases general commodity price levels, but only slightly. When either real interest rates decrease or the dollar appreciates, oil prices increase and exhibit "overshooting" behavior (Akram 2009). Cologni and Manera (2008) show shocks to the price of oil cause an instantaneous but temporary effect on price levels in most G-7 countries.

Jones and Kaul (1996) find US and Canadian stock prices react to shocks in oil prices through their effects on current and future cash flows. Oil prices have played a larger role than interest rates on the forecast error variances of stock returns since the

middle of the 1980s (Sadorsky 1999). Oil prices increase in the months following an interest rate shock, and there are smaller, less sustained increases in oil price in response to shocks in industrial production and stock returns. On the other hand, a shock in oil prices leads to an immediate decrease in stock returns that appears to stabilize after one quarter. Huang, Masulis, and Stoll (1996), however, find no relationship between returns for oil futures and returns in the equity markets except for oil companies.

Oil Prices and the Energy Industry

Activities of the petroleum industry are commonly classified as upstream, midstream, or downstream. Upstream activities are concerned primarily with exploration, drilling, and extraction; midstream activities involve the transportation and storage of raw products; and downstream activities are associated with refiners' production of consumable products. While the literature on the dynamics of the energy industry as a whole is limited, substantial research has been performed on specific sectors. Ederington et al. (2011) provide discussions of these topics.

Upstream

Upstream activities revolve around exploration and discovery. Oil production demonstrates unit root behavior for many producing countries (Maslyk and Smyth 2009), so production shocks, therefore, likely have a non-transitory effect on economic output. Güntner (2014) finds oil production responds to speculative demand shocks at a lag but has no statistically significant responses to oil demand shocks. The number of drilling sites is inelastic with respect to oil prices but more elastic than the volume of

reserves (Smith and Lee 2017), evident during the 2014 decrease in oil prices as reductions in reserve additions and production were less than those in drilling activity.

Rig count is a common metric of this sector's level of activity (Cheng 1998; Dahl and Duggan 1998) and commonly used as a proxy for oil production. Chen and Linn (2017) find a lagged relationship between changes in field production of oil and natural gas and changes in futures prices for oil and gas. Guerra (2008) uses rig counts as a proxy for investment and finds at least 40 percent of rig activity is explained by oil price. Oil prices show only a small and temporary response to changes in rig count. Mohn (2008) examines the relationship between oil prices and production though changes in drilling efforts and efficiency. He finds a long-run relationship between oil prices, drilling efforts, success, and discovery size. Cheng (1998) finds no Granger causality between wells drilled and oil price. This result holds even when rig count is substituted for the number of wells drilled. In the long-run, however, prices are found to Granger-cause wells drilled. Olatubi and No (2003) find both price shocks and shocks in price volatility increase exploration activity in the short run. Dahl and Duggan's (1998) survey of price elasticities reveals rig activity responds to prices. Some of the responses are slowly realized, and all but one of the surveyed long-run price elasticities are greater than one. They suggest, however, the omission of drilling costs and the effect of geography may bias the estimated coefficients. Ringlund, Rosendahl, and Skjerpen (2008) find rig activity increases following an increase in oil prices. The response of the rig count to the price increase is noticeable after three months in the US but at later times in less developed countries. They estimate the long-run elasticity of rig count with

respect to oil prices is 1.3. Mohn and Osmundsen (2008) study drilling activity on the Norwegian Continental Shelf and determine drilling activity is best proxied by the number of wells when comparing alternative measures, such as number of drilling days, and expenditures on exploration. The long-run elasticity of drilling activity with respect to price is between 0.20 and 0.41. There is evidence of a long-term effect on drilling of oil price, but over 80 percent of the short-term deviations from equilibrium are adjusted within one year.

Midstream

Compared to the upstream and downstream, the midstream is researched less frequently. One component of midstream operations is storage. The economic theory of storage has its origins in Working (1949), who suggests there is an implicit price for storage necessary to understand futures prices. Pindyck (2001, p. 2) shows "…inventories play a crucial role in price formation …" because they "…reduce costs of changing production." Inventories also dampen price volatility (Du, Yu, and Hayes 2011). Unalmis, Unalmis, and Unsal (2012) find the impact of supply shocks on oil prices is overstated when the model does not account for storage. Ye, Zyren, and Shore (2002) explore the performance of forecasts for crude prices using OECD petroleum inventories. Their best forecasting model includes current and lagged deviations from normal OECD inventory levels. The ratio of inventory to sales decreases following a shock in oil price (Herrera 2018).

Transportation, too, is an important part of midstream operations. Interruptions in the flow of petroleum products from upstream to downstream or from downstream to

retailers affect prices. In September 2016 gasoline futures, for example, increased when a pipeline leak was detected in Alabama (Friedman 2016). Olsen (2011) finds evidence of cointegrating relationships between North American natural gas markets, suggesting regional deviations from the "one price" are due to transportation and transaction costs. Olsen, Mjelde, and Bessler (2015) cite bottlenecks as a reason markets separated by greater distances are less integrated. Mu (2007) finds the volatility of natural gas prices increases the day storage information is released. Pipeline capacity constraints also increase wholesale natural gas prices for regional markets. For the Florida Gas Transmission Company between 2006 and 2011, pipeline capacity utilization of 95 percent or more raised average regional prices over 11 percent (Avalos, Fitzgerald, and Rucker 2016).

Downstream

Gabel (1979), using a simultaneous equation model of the US refining industry, finds increased capacity utilization causes an increase in the price markup. There is a negative relationship between the number of refineries and capacity utilization. Kaufmann et al. (2008) find a negative relationship between refinery utilization rates and oil prices. Reducing refinery activity by one percentage point would increase oil prices four percent. In addition, because light, sweet crudes command higher prices and revenues, refinery utilization rates may also affect the spread between prices for heavy and light crude oil. Dèes et al. (2008), on the other hand, show the increase in crude oil prices from 2004 to 2006 was caused by conditions in the futures markets and a decrease in

refinery capacity utilization, but they conclude an increase in refinery capacity would not decrease crude oil prices.

Adams and Griffin (1972) formulate a linear programming model of the refining industry. At 88.1 percent utilization, output prices decrease because marginal costs are low, leading to price wars. Policies affecting oil prices alter the prices for end products.

Chesnes (2009) finds the optimal utilization rate for refineries increases early in the year and peaks in late summer. Shocks in oil prices are passed through to the consumers of end products, but price increases in gasoline are less than the increase in refiners' costs. A 20 percent shock in crude price decreases consumer surplus by 58 percent, refiners' profits by 37 percent, and total welfare by 45 percent in the year following the shock. A late summer hurricane similar in magnitude to Katrina or Ike is hypothesized to cause a 25 percent decrease in refinery capacity and as a result almost a 70 percent decrease in consumer surplus, a 15 percent increase in profits, and a decrease in total welfare of 11 percent. Unplanned outages decrease capacity utilization (Chesnes 2015). Planned outages, on the other hand, tend to occur when operating margins are lowest. Outages, in general, tend to increase the price of refined products, but the amount of the increase depends on the available capacity.

While the dynamics of gasoline prices are frequently investigated, sales of gasoline are studied less frequently. Gasoline prices demonstrate asymmetry in response to increases and decreases in oil prices (Radchenko 2005). Kuper (2012) decomposes the asymmetry, finding higher marketing and storage costs cause prices to increase fast than they decrease. Similar to oil prices, gasoline prices are explained in the long-run

substantially (68 percent of the variation) by oil and aggregate demand shocks, while in the short-run, gasoline prices are primarily explained by gasoline demand shocks (Kang, Perez de Gracia, and Ratti 2018).

The FAVAR Model

Bernanke, Boivin, and Eliasz (2005), following Stock and Watson (2002a, c), introduce the FAVAR. Researchers often have substantial information at their disposal, but a lot of this information is not included in typical VAR analyses. The VAR's inability to incorporate the information from many variables results in "sparse information," necessitating the researcher to "…[take] a stand on specific observable measures corresponding precisely to some theoretical constructs" (Bernanke, Boivin, and Eliasz 2005, p. 389). The FAVAR approach combines the strengths of traditional VAR analysis – the ability to model dynamic relationships among variables through impulse response functions (IRFs) and forecast error variance decompositions (FEVDs) – with the ability to analyze many series simultaneously. In their empirical application, Bernanke, Boivin, and Eliasz (2005), for example, estimate the relationships between one policy variable, the federal funds rate, and 120 informational series, including measures of output and income, employment, consumption, housing activity, inventories, stock indices, exchange, interest, and inflation rates, and money supply.

The FAVAR is built on statistical factor analysis, a technique commonly used for dimension reduction (Rencher and Christensen, 2012). Statistical factor analysis emerged as a cross-sectional data analysis technique concerned with choosing an appropriate number of latent factors k that adequately captures the features of n variables

and then estimating both the factors f_j and their contributions λ_{ij} to each of the observed variables y_i. Dimension reduction is achieved by k being less than n. If each of the observed variables y_i has mean μ_i, then the orthogonal factor model is $y - \mu = \Lambda f + e$, where y is a vector of the original, observed series, μ is a vector of the means of each variable, f is a vector of the factors, and e is a vector of residuals. The conformable matrix Λ contains factor loadings, the linear transformations of the factors into the observed variables. Many techniques are available for estimating the factors and their loadings. As an application of statistical factor analysis, the FAVAR model tries to capture the features of many time series.

Zagaglia (2010) employs the FAVAR to investigate the term structure of oil futures with many informational series. His FAVAR includes 230 series and improves out-of-sample, one- and three-step ahead forecasts of oil futures prices relative to a traditional VAR, a VAR of factors only, and a random walk. Ipatova (2014), using an updated version of Zagalia's (2010) dataset, shows forecasts of oil futures are improved using the FAVAR approach. Four oil futures prices are modeled in addition to 301 informational series. An, Jin, and Ren (2014) use the FAVAR to study the asymmetry in the responses of macroeconomic variables to oil price shocks. Using quarterly data on 114 macroeconomic series, they find rising oil prices decrease, among others, output, saving, payrolls, and house prices. The asymmetry in the responses to positive versus negative oil price changes is larger in the short run compared to the long run. Aastveit (2014) also studies oil price shocks with the FAVAR with 112 informational series and concludes "…it is important to consider the causes behind the movements of oil price"

(Aastveit 2014, p. 278). Binder, Pourahmadi, and Mjelde (forthcoming) show different methods for factor extraction affects the ability of FAVAR models to forecast oil prices. Other studies employing the FAVAR but unrelated to oil prices include Forni and Gambetti (2010), Bahadir and Lastrapes (2015), Lutz (2015), Ratti and Vespignani (2015), Duangnate (2015), and Binder (2016).

The Time-Varying Parameter FAVAR Model

The FAVAR model of Bernanke, Boivin, and Eliasz (2005) assumes a *k*-vector of latent factors f_t captures the movements of an *n*-vector of informational series x_t. The latent factors are assumed to evolve dynamically with an *m*-vector of observational variables y_t in the VAR framework

$$\Phi(L)\begin{bmatrix} f_t \\ y_t \end{bmatrix} = v_t, \tag{3.1}$$

where $\Phi(L)$ is a lag polynomial matrix and v_t is a vector of zero-mean innovations such that $E(v_t v_t') = \Omega$. The informational series, latent factors, and vector of observational variables are related by

$$x_t = \Lambda_f f_t + \Lambda_y y_t + e_t, \tag{3.2}$$

where Λ_f and Λ_y are $n \times k$ and $n \times m$ matrices of coefficients and e_t is assumed to be an uncorrelated vector of errors with diagonal covariance matrix *R*.

Mumtaz, Zabczyk, and Ellis (2011) build on the Bernanke, Boivin, and Eliasz (2005) FAVAR framework with Primiceri's (2005) time-varying parameter (TVP) VAR. The TVP-FAVAR with *l* lags can be written in state-space form with observation equation

(3.3) $$\begin{bmatrix} x_{t,n\times 1} \\ y_{t,m\times 1} \end{bmatrix} = \begin{bmatrix} \Lambda_{f,n\times k} & \Lambda_{y,n\times m} \\ 0_{m\times k} & I_{m\times m} \end{bmatrix} \begin{bmatrix} f_{t,k\times 1} \\ y_{t,m\times 1} \end{bmatrix} + \begin{bmatrix} e_{t,n\times 1} \\ 0_{m\times 1} \end{bmatrix},$$

where I is the identity matrix and all other variables are as previously defined. The vector of observational series includes both the Cushing West Texas Intermediate crude oil spot price (oil prices) and the federal funds rate (interest rate). Here, the n-vector of informational series x_t includes many series describing each sector of the petroleum industry. The informational series also include many macroeconomic series, representing current economic activity. The inclusion of many industry-specific time series allows for the simultaneous estimation of these series' responsiveness to changes in oil price and the federal funds rate. The latent factors f_t span the space of the entire set of informational series, meaning the information contained in each of the series, both the industry-specific and macroeconomic, contributes to the values of each factor.

For the transition equation, the dynamics of the latent factors and observational series are jointly specified as a VAR

(3.4) $$\begin{bmatrix} f_{t,k\times 1} \\ y_{t,m\times 1} \end{bmatrix} = \Phi_t \begin{bmatrix} f_{t-1,k\times 1} \\ y_{t-1,m\times 1} \end{bmatrix} + v_t,$$

where Φ_t is a conformable matrix of time-varying coefficients and v_t is a vector of residuals. On occasion, the VAR will also be written in the shorthand notation

(3.5) $$F_t = \Phi_t F_{t-1} + v_t,$$

where $F_t = \begin{bmatrix} f_t' & y_t' \end{bmatrix}'$. While this notation implies only one lag of the state vector is included in the VAR, additional lags can be incorporated by considering equation (3.5)

to represent the companion or VAR(1) form of a multiple-lag VAR model (Lutkepohl 2005).

It is assumed the VAR coefficients follow a random walk

$$vec(\Phi_{l,t}) = vec(\Phi_{l,t-1}) + \eta_t, \tag{3.6}$$

where $vec(\cdot)$ is the vectorization operator, $\Phi_{l,t}$, is the lth member of the lag polynomial, and η_t is a vector of white noise with covariance matrix Q.

The covariance matrix $E(v_t v_t')$ also evolves over time and can be decomposed

$$E(v_t v_t') = \Omega_t = A_t^{-1} H_t H_t' (A_t^{-1})'. \tag{3.7}$$

The matrix H_t is a diagonal matrix of the standard deviations of the innovations for each equation of the VAR. Let h_t denote the vector of the diagonal elements of H_t. Each standard deviation is assumed to evolve independently according to the stochastic volatility model of Kim, Shephard, and Chib (1998) such that

$$\ln(h_t) = \ln(h_{t-1}) + \xi_t \tag{3.8}$$

where ξ_t is a vector of independently distributed white noise distributed according to diagonal covariance matrix G.

The matrix A_t describes the contemporaneous relationships between the variables of the VAR system such that

$$A_t(F_t - \Phi_t F_{t-1}) = H_t w_t, \tag{3.9}$$

where

$$\Phi_t = \begin{bmatrix} \Phi_{1,t} & \Phi_{2,t} & \dots & \Phi_{l-1,t} & \Phi_{l,t} \\ I & 0 & \dots & 0 & 0 \\ 0 & I & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & I & 0 \end{bmatrix},$$

each $\Phi_{l,t}$ is the coefficient matrix at time *t* associated with the *l*th lag of the VAR, and w_t is a vector of independent and identically distributed white noise with unit variance. The matrix A_t is assumed to have a lower triangular structure scheme. Let a_t denote a vector of the unrestricted values of A_t, the lower-triangular, non-diagonal members only, ordered by row and column such that

$$A_t = \begin{bmatrix} 1 & 0 & 0 & 0 & \dots \\ A_{2,1,t} & 1 & 0 & 0 & \dots \\ A_{3,1,t} & A_{3,2,t} & 1 & 0 & \dots \\ A_{4,1,t} & A_{4,2,t} & A_{4,3,t} & 1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

corresponds to the vector $a_t = \begin{bmatrix} A_{2,1,t} & A_{3,1,t} & A_{3,2,t} & A_{4,1,t} & \dots \end{bmatrix}'$. The elements of this vector are assumed to evolve according to a random walk

$$a_t = a_{t-1} + \varepsilon_t, \tag{3.10}$$

where ε_t is a vector of Normally distributed white noise with covariance matrix *S*.

Bernanke, Boivin, and Eliasz (2005) employ a recursive identification scheme using the Cholesky decomposition to perform innovation accounting, ordering the monetary policy variable last. In this dissertation, Uhlig's (2005) sign restriction technique, in conjunction with the elasticity bound of Kilian and Murphy (2010), is used

to set identify the structural responses, so the ordering of the series does not matter for identification of the structural shocks. This is discussed in detail later.

Estimation

The TVP-FAVAR model is complex with numerous parameters to estimate. There are hundreds of informational series, leading to many equations of factor loadings; there are also time-varying coefficients for the VAR, the time-varying volatilities, and time-varying contemporaneous relationships between the variables in the VAR. In addition, the model is nonlinear with respect to the log volatility of the transition equation. As a result, estimators requiring optimization of global extrema are likely to be computationally infeasible.

Bernanke, Boivin, and Eliasz (2005) propose two estimation methods: a two-step estimator using principal components that first estimates the factors and then their loadings and VAR coefficients. Their second method, a one-step method, uses Bayesian methods to estimate the posterior distribution of all parameters jointly. Because they find little difference in the numerical results, the simplicity of the two-step method is appealing. The two-step estimator, however, does not provide a ready solution for the estimation of the TVP-FAVAR. Korobilis (2012) applies a two-step estimation with time-varying parameters by first estimating and identifying the factors using the principal components technique and then applying the framework of the Bayesian TVP-VAR (Primiceri 2005, Cogley and Sargent 2005) extended to the FAVAR. Mumtaz, Zabczyk, and Ellis (2011), however, use Bayesian methods to estimate all the model's

parameters, obtaining even the factors through Bayesian methods. Bayesian methods are pursued here to estimate jointly the factors and the parameters of the TVP-FAVAR.

Overview of Bayesian Methods. The Bayesian paradigm is summarized by Gelman et al. (2013). Bayesian statistical estimation and inference is built on probability statements conditioned on the observed data. Gelman et al. (2013, p. 6) explain:

> It is at the fundamental level of conditioning on observed data that Bayesian inference departs from the approach to statistical inference described in many textbooks, which is based on a retrospective evaluation of the procedure used to estimate [the unknown parameters] over the distribution of possible [observed values] conditional on the true unknown value of [the parameters].

Bayesian inference treats unknown parameters as random variables and attempts to estimate the probability distribution describing the potential values of the parameters (Gelman et al. 2013; Hoff 2009). The probability distribution depends on the likelihood of the data and the investigator's prior beliefs regarding the values of the parameters being estimated. In Bayesian statistics, observed data are treated as given and are the only values known to the researcher. Samples from the population that could have been drawn or would be drawn in repeated sampling are not of interest. In contrast, many frequentist estimators are evaluated by considering their performance under other potential draws of the observed data. A common performance metric of frequentist estimators is bias; an estimator is said to be unbiased if, in resampling from the population, the expectation of its sample mean is the true value of the parameter being estimated. While the results of Bayesian inferences are often numerically similar to those of their frequentist counterparts, Bayesian analysis has an advantage in its ability to handle large, complex models such as the TVP-FAVAR.

In the TVP-FAVAR model, the observed data include informational series x_t and observational series y_t. The unknown parameters of interest are:

- f_t, the latent factors for time $t = 1,\ 2,\ ...,\ T$;
- Λ_f and Λ_y, the matrices of coefficients in the measurement equation (3.3)
- R, the matrix of the variances of the residuals of the measurement equation;
- Φ_t, the matrix of the time-varying coefficients of the VAR in the transition equation (3.4) for time $t = 1,\ 2,\ ...,\ T$;
- H_t, the diagonal matrix of the time-varying volatility of the VAR residuals from the transition equation for time $t = 1,\ 2,\ ...,\ T$;
- A_t, the lower-triangular matrix of the time-varying contemporaneous relationships of the variables in the VAR in the transition equation for time $t = 1,\ 2,\ ...,\ T$;
- Q, the covariance matrix of the innovations to the random walk governing the evolution of the coefficients in Φ_t, as defined in equation (3.6);
- G, the covariance matrix of the innovations to the random walk governing the evolution of the coefficients in H_t as defined in equation (3.8); and
- S, the covariance matrix of the innovations to the random walk governing the evolution of the coefficients in A_t as defined in equation (3.10).

Following Kim and Nelson (1999, p. 197), let the "~" notation denotes the entire set of values of the time-varying random variables across all periods, for example, $\tilde{x} = \{x_1, x_2, \ldots, x_t\}$. Define $p(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G)$ as the prior distribution of the unknown parameters. The prior distribution is a probabilistic summary of the investigator's beliefs regarding the unknown parameters. The likelihood of the data $p(\tilde{x}, \tilde{y} \mid \tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G)$ is the probability of the observed the data conditional on the parameters. Probability statements on the parameters conditional on the observed data are made using Bayes' Theorem to obtain

(3.11)

$$
\begin{aligned}
p(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G \mid \tilde{x}, \tilde{y}) & \\
&= \frac{p(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G)\, p(\tilde{x}, \tilde{y} \mid \tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G)}{p(\tilde{x}, \tilde{y})} \\
&\propto p(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G)\, p(\tilde{x}, \tilde{y} \mid \tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G),
\end{aligned}
$$

where $p(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G \mid \tilde{x}, \tilde{y})$ is the posterior distribution, the probability distribution of the parameters conditional on the observed data. By the properties of probability, the left-hand side of the relationship in equation (3.11) can be rewritten (Kim and Nelson 1999, p. 197) as

$$
\begin{aligned}
&= p(\Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G \mid \tilde{f}, \tilde{x}, \tilde{y})\, p(\tilde{f} \mid \tilde{x}, \tilde{y}) \\
&= p(\tilde{\Phi}, \tilde{H}, \tilde{A}, S, Q, G \mid \tilde{f}, \tilde{x}, \tilde{y})\, p(\Lambda, R \mid \tilde{f}, \tilde{x}, \tilde{y})\, p(\tilde{f} \mid \tilde{x}, \tilde{y}) \\
&= p(S \mid \tilde{\Phi}, \tilde{H}, \tilde{A}, \tilde{f}, \tilde{x}, \tilde{y})\, p(Q \mid \tilde{\Phi}, \tilde{H}, \tilde{A}, \tilde{f}, \tilde{x}, \tilde{y})\, p(G \mid \tilde{\Phi}, \tilde{H}, \tilde{A}, \tilde{f}, \tilde{x}, \tilde{y}) \times \\
&\quad p(\tilde{\Phi}, \tilde{H}, \tilde{A}, \mid \tilde{f}, \tilde{x}, \tilde{y})\, p(\Lambda, R \mid \tilde{f}, \tilde{x}, \tilde{y})\, p(\tilde{f} \mid \tilde{x}, \tilde{y}).
\end{aligned}
$$

This equality simplifies the posterior into several conditionally independent components, leading to a simplified estimation strategy.

1) Conditional on the observational series, the informational series, and the other parameters of the model, draw the latent factors.
2) Conditional on the latent factors and the observational series, draw the factor loadings and the residual covariance of the observation equation from equation (3.3).
3) Conditional on the latent factors and the observational series, draw the parameters of the transition equation, equations (3.4), (3.6), (3.8), and (3.10):
    a. the time-varying VAR coefficients and the covariance matrix governing their evolution,
    b. the time-varying volatilities and the variance parameters governing their evolution, and
    c. and the time-varying contemporaneous relations and the covariance matrix governing their evolution.

A detailed description of each step of the iteration is now provided.

Step 1. Draw the Latent Factors. Conditional on all the parameters of the model, the estimation focuses on equations (3.3) and (3.4). Assuming e_t is Normally distributed with covariance matrix R, v_t is Normally distributed with covariance matrix Ω_t, and e_t and v_t are independent, these equations constitute a linear, Gaussian state space model. As such, the Kalman Filter, "…a recursive procedure for computing the optimal estimate of the unobserved-state vector…" (Kim and Nelson 1999, p. 22), can be applied to obtain the conditional expectation and covariance of the latent state vector F_t at each

period. Applied to this state-space model, the Kalman Filter equations for each period are

$$
\begin{aligned}
F_{t|t-1} &= \Phi_{t-1} F_{t-1} \\
P^F_{t|t-1} &= \Phi_{t-1} P^F_{t-1|t-1} \Phi_{t-1}{}' + \Omega_{t-1} \\
u_{t|t-1} &= \begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} - \begin{bmatrix} \Lambda_f & \Lambda_y \\ 0 & I \end{bmatrix} F_{t|t-1} \\
W_{t|t-1} &= \begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} P^F_{t|t-1} \begin{bmatrix} x'_{t-1} & y'_{t-1} \end{bmatrix} + R \\
F_{t|t} &= F_{t|t-1} + K_t u_{t|t-1} \\
P_{t|t} &= P^F_{t|t-1} - K_t \begin{bmatrix} \Lambda_f & \Lambda_y \\ 0 & I \end{bmatrix} P^F_{t|t-1},
\end{aligned}
\tag{3.12}
$$

where, following the notation of Kim and Nelson (1999):

- $F_{t|t-1}$ is the conditional expectation of the latent state vector F_t at time t as in equation (3.5) given information up to $t-1$;
- $P^F_{t|t-1}$ is the conditional covariance of the latent state vector at time t given information up to $t-1$;
- $u_{t|t-1}$ is the conditional forecast error of the observation equation at time t given information up to $t-1$;
- $W_{t|t-1}$ is the conditional covariance of the forecast error of the observation equation at time t given information up to $t-1$; and
- $K_t = P^F_{t|t-1} \begin{bmatrix} \Lambda_f & \Lambda_y \\ 0 & I \end{bmatrix}' W^{-1}_{t|t-1}$ is the Kalman gain at time t, a weighting matrix used to update the expectation and covariance of the state vector.

To use the Kalman Filter technique, initial values for $F_{0|0}$ and $P_{0|0}^{F}$ must be provided. Because the coefficients of the VAR in the transition equation follow a random walk and are unrestricted, the latent factor vectors are not guaranteed to be stationary, so the unconditional mean and covariance of the factors do not exist. The recursions, therefore, are initialized with $F_{0|0} = \left[\underbrace{0 \quad \ldots \quad 0}_{k} \quad y_{1,0} \quad y_{2,0} \right]'$, where $y_{1,0}$ is the previous known value of the first observational variable (oil prices), $y_{2,0}$ is the previous known value of the second observational series (federal funds rate), and the sequence of *k* zeros correspond to the starting values for the *k* factors. The prediction covariance $P_{0|0}^{F}$ is initialized as a diagonal matrix with one on each element of the diagonal corresponding to the positions of the factors to be estimated and zeroes in the positions corresponding to the known values of the observed series.

Once the Kalman Filter iterations are complete, backward recursions based on Carter and Kohn (1994) as outlined in Kim and Nelson (1999) are used to obtain samples from the full conditional posterior of the factors. To begin, draw a value of the state vector at time *T*, the final observation, from $N(F_{t|t}, P_{t|t}^{F})$, where *N* denotes the Normal distribution with mean $F_{t|t}$ and covariance $P_{t|t}^{F}$. The updating steps for times $T-1$ to 1 are

$$
\begin{aligned}
F_{t|t,F_{t+1}^{*}} &= F_{t|t} + P_{t|t}^{F}\Phi_{t}'(\Phi_{t}P_{t|t}^{F}\Phi_{t}' + \Omega_{t})^{-1}(F_{t+1}^{*} - \Phi_{t}F_{t|t}) \\
P_{t|t,F_{t+1}^{*}}^{F} &= P_{t|t}^{F} - P_{t|t}^{F}\Phi_{t}'(\Phi_{t}P_{t|t}^{F}\Phi_{t}' + \Omega_{t})^{-1}\Phi_{t}P_{t|t}^{F},
\end{aligned}
\tag{3.13}
$$

where $F_{t|t,F_{t+1}^*}$ is the conditional expectation at time *t* of the state vector given information up to time *t* and F_{t+1}^*, the draw of the state vector at time $t+1$, and $P_{t|t,F_{t+1}^*}$ is the conditional covariance at time *t* of the state vector given information up to time *t*. When the updated conditional mean and covariance for time *t* is obtained, a draw is made from $N(F_{t|t,F_{t+1}^*}, P_{t|t,F_{t+1}^*}^F)$ to obtain F_t^*, this iteration's draw for the state vector at time *t*. Repeat the updating steps for each period until the first observation is reached.

In applications where the number of lags in the transition equation, exceeds one, a modification is needed to the backwards recursions because of the singularity of Ω_t. In the backwards recursions, all instances of the matrices Φ_t and Ω_t are truncated after $m+k$ rows. Kim and Nelson (1999) provide more details on this procedure.

For the TVP-FAVAR estimated by Bayesian methods, there are two components to identification. The first is the identification of the state space model to prevent rotation of the factors; the second is the identification of structural innovations for innovation accounting (discussed later). The state space model specified here is unidentified without any restrictions on the coefficients of the observation or transition equations. Bernanke, Boivin, and Eliasz (2005) show that in their one-step, Bayesian estimation method, identification is achieved using information from both the observation and transition equations. To prevent rotation of the factors, restrictions must be placed on the coefficient matrices of either the observation or transition equation. Following Bernanke, Boivin, and Eliasz (2005) and Mumtaz, Zabczyk, and Ellis (2011), the first *k* rows of the matrix of coefficients from equation (3.3) are set to zero. In the

first row, the coefficient corresponding to the first latent factor is set to one; in the second row, the coefficient corresponding to the second latent factor is set to one; similar restrictions are placed through the *k*th row of the coefficient matrix.

Step 2. Draw Observation Equation Parameters. Conditional on the factors, the observation equation of the state-space model in equation (3.3) is independent of the rest of the model. This is a system of regression equations. The regressors for each equation are the latent factors generated previously. In cases where the number of lags in the transition equation exceeds one, only the values of F_t and not F_{t-1}, F_{t-2}, ..., F_{t-l} are included as regressors in the observation equation. It was previously assumed the vector e_t is Normally distributed with covariance matrix *R*. Now, it is further assumed *R* is diagonal. Because of the independence of the innovations across equations, each row of the observation equation can be treated as a separate regression model and estimated separately. A conjugate prior for the single regression model is the multivariate Normal-Inverse Gamma prior such that the prior of *i*th row of the matrix of coefficients is $p\left(\left[\Lambda_f \quad \Lambda_y\right]_i\right) = N(0, I)$. For $R_{i,i}$, the *i*th diagonal element of *R*, the prior is Inverse Gamma $p(R_{i,i}) = IG(0.01, 0.01)$, denoting the shape and rate parameters. The choice of the identity covariance matrix for the coefficients matches that of Bernanke, Boivin, and Eliasz (2005) and Amir-Ahmadi and Uhlig (2009). Experimentation with Jeffrey's noninformative prior and a uniform prior resulted in difficulties with numerical stability of the Markov Chain Monte Carlo (MCMC) algorithm and caused explosive behavior in the parameters of *R* and Ω_t. The conditional posterior for the coefficients is Gaussian

$N\left((I+F'F)/R_{i,i})^{-1}(F'X_i/R_{i,i}),\ (I+F'F)/R_{i,i})^{-1}\right),$ where F is the matrix of the latent state vectors, X_i represents the column vector of the values of the ith dependent variable in the matrix equation (Hoff 2009). The conditional posterior of the residual variance is Inverse Gamma $IG\left(0.01+0.5\times T, 0.01+0.5\times \text{SSR}_i^R\right),$ where T is the number of observations and SSR_i^R is the sum of squared residuals for the ith regression equation (Hoff 2009).

Step 3a. Draw the Transition Equation Time-Varying VAR Coefficients. Conditional on the factors and $\Omega_t,$ the transitional equation from (3.4) becomes

$$F_t = \begin{bmatrix} F'_{t-1} & 0 & \dots & 0 \\ 0 & F'_{t-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & F'_{t-1} \end{bmatrix} vec(\Phi_t) + v_t,$$

where the time-varying VAR coefficients evolve $vec(\Phi_t) = vec(\Phi_{t-1}) + \eta_t.$ Previously, it was assumed v_t is Normally distributed with covariance matrix $\Omega_t.$ It is additionally assumed η_t is Normally distributed with covariance matrix Q and that v_t and η_t are independent. As a result, these equations also constitute a linear, Gaussian state space model estimable by Kalman Filter as discussed previously. The Kalman Filter recursions are started with a zero vector except for the coefficient on each variable's first lag which is set to 0.9. The starting covariance is specified as a diagonal matrix with 0.0001 for each diagonal element. The backward recursions of Carter and Kohn (1994) are then applied to obtain draws for the coefficients.

The covariance matrix Q is then updated. Cogley and Sarget (2005) and Primiceri (2005), along with the experience garnered during the estimation of this model, suggest the choice of prior for Q is key to maintaining non-explosive behavior of the VAR coefficients as they evolve in the random walk. The literature proposes a relatively tighter prior on this covariance matrix, so the prior is specified as the Inverse Wishart distribution $p(Q) = IW(\delta, 0.0001 \times \delta \times I)$, where I is an identity matrix of the same dimensions as Q and the constant δ is set at one more than the number of time-varying coefficients in the VAR (the dimension of Q plus one). For the Inverse Wishart distribution, having degrees of freedom exceeding the dimension of the matrix guarantees the matrix is positive definite (Hoff 2009). Because the Inverse Wishart is naturally conjugate to the Normal density, the conditional posterior is also Inverse Wishart $IW\left(T-1+\delta, (0.0001 \times 40 \times I + U_Q)^{-1}\right)$, where

$$U_Q = \sum_t \left[vec(\Phi_t) - vec(\Phi_{t-1})\right]\left[vec(\Phi_t) - vec(\Phi_{t-1})\right]'.$$

Step 3b. Draw the Time-Varying Volatilities of the Transition Equation. Conditional on the factors, the time-varying coefficients, and the time-varying contemporaneous relationships, values for the time-varying volatilities can be drawn. Following Primiceri (2005) and Del Negro and Primiceri (2015), the transition equation as given in equation (3.9) is a nonlinear matrix equation. On the right side of the equation is the matrix of standard deviations multiplied by white noise. This equation is linearized by first squaring both sides. By squaring both sides, the diagonal matrix H_t of standard deviations becomes a diagonal matrix of variances. The natural logarithm is taken on

both sides to complete the linearization, forcing the product on the right side of the equation to become the sum of the natural logarithms of each component. The linearized equation is

$$\hat{f}_t = 2\ln h_t + e_t, \tag{3.14}$$

where $\hat{f}_t = \ln\left([A_t(F_t - \Phi_t F_{t-1})]^2 + c\right)$, $e_{it} = \ln w_{it}^2$, and *i* indexes the rows of the vector $\hat{f}_t$. The constant *c* is used for numerical stability when taking the natural logarithm of small values and is set to an arbitrarily small value, here 0.00001. In equation (3.8), it was assumed the log volatility of the transition equation evolved as a random walk. Together, equations (3.8) and (3.14), yield the state space model

$$\begin{aligned} \hat{f}_t &= 2\ln h_t + e_t \\ \ln h_t &= \ln h_{t-1} + \xi_t, \end{aligned} \tag{3.15}$$

where ξ_t is Normally distributed white noise with diagonal covariance matrix *G*. The innovations of the two equations are assumed independent of each other, but the innovations of the observation equation are distributed according to the $\log \chi^2$ distribution with one degree of freedom. As a result, this state space model is linear but not Gaussian. Following Del Negro and Primiceri (2015) and Kim, Shephard, and Chib (1998), the density of the $\ln \chi^2$ distribution can be approximated by a mixture of Gaussian distributions described in Table 4 of Kim, Shephard, and Chib (1998) because a linear combination of Normal distributions yields a Normal distribution. The moments of each component distribution are selected such that the moments of the mixture of Normal distribution are close to those of the $\ln \chi^2$ distribution. Let s_{it} indicate which of

the component distributions is applied to $\hat{f}_{i,t}$. The $s_{i,t}$ can be sampled from the multinomial distribution

$$\Pr(s_{i,t} = j) \propto q_j \phi\left(A_t(F_t - \Phi_t F_{t-1}) \mid 2\ln h_{i,t} + \mu_j - 1.2704, \sigma_j^2\right), \tag{3.16}$$

where ϕ is the Normal density with mean $2\ln h_{i,t} + \mu_j - 1.2704$ and variance σ_j^2 and all constants are as given in Table 4 of Kim, Shephard, and Chib (1998) .

Upon sampling $s_{i,t}$ for all variables in all periods, the covariance matrix of e_t at each period is constructed. The variance of $e_{i,t}$ is the variance of the $s_{i,t}$ component of the mixture of Normal distributions. The system given in equation (3.15) is now a linear, approximately Gaussian state space model, and the Kalman Filter and backward recursions can be applied to sample values for $\ln h_{i,t}$. The Kalman Filter is initialized with a zero vector and identity covariance matrix.

Del Negro and Primiceri (2015), a corrigendum to Primiceri's (2005) original TVP-VAR, provide an additional step that corrects for the fact the likelihood of the $\ln \chi^2$ distribution is approximated by the mixture of Normal distributions proposed by Kim, Shephard, and Chib (1998). Upon obtaining sampled values for the time-varying volatilities using the Carter and Kohn (1994) backwards recursions, a Metropolis-Hastings step is used to accept or reject the drawn values. The proposed value is accepted with probability

$$\alpha = \frac{\left[\prod_t \phi(v_t \mid 0, H_t^{**})\right]\left[\prod_t \prod_i \tilde{\phi}(\hat{f}_{i,t} - 2\ln h_{i,t}^{*})\right]}{\left[\prod_t \phi(v_t \mid 0, H_t^{*})\right]\left[\prod_t \prod_i \tilde{\phi}(\hat{f}_{i,t} - 2\ln h_{i,t}^{**})\right]}, \tag{3.17}$$

where $\tilde{\phi}$ is the density of the mixture of Normal distributions approximation to the $\ln \chi^2$ distribution $\tilde{\phi}(\cdot) = \sum_j \phi(\cdot|\mu_j - 1.2704, \sigma_j^2)$, the "**" denotes the previous value of the Markov Chain, and the "*" denotes the value proposed in the current iteration of the chain.

An additional benefit of the $\hat{f}_t$ transformation and the diagonal assumption for *G* is that each row of the observation equation in (3.15) is independent of the others. This allows for an improved single-component Metropolis-Hastings update. The single-component strategy increases the sampler's efficiency because the acceptance or rejection of one row's proposed values is determined without considering the proposed values for the other rows. In this application, it allows the sampler to accept and update the Markov Chain more frequently than if all rows of the vector h_t are accepted or rejected simultaneously. To perform the single-component update, the Kalman Filter, backwards recursions, and acceptance/rejection steps are performed one row at a time.

The diagonal covariance matrix *G* is then updated. Following Cogley and Sarget (2005), the conjugate Inverse Gamma distribution is selected as the prior for each diagonal element of for each $p(G_{ii}) = IG\left(0.5 \times 0.0001, 0.5\right)$. This yields an Inverse Gamma conditional posterior $IG\left(0.5 \times (0.0001 + T), 0.5 \times (1 + \mathrm{SSR}_i^G)\right)$, where SSR_i^G is the sum of squared residuals ξ_t from the random walk of the *i*th log volatility.

Step 3c. Draw the Time-Varying Contemporaneous Relations of the Transition Equation. Conditional on the factors, the time-varying coefficients, and the time-varying

innovation volatilities, values for the time-varying contemporaneous relationships are drawn. By the lower-triangular assumption for A_t, equation (3.9) can be rewritten as

$$F_t - \Phi_t F_{t-1} = \begin{bmatrix} 0 & 0 & \dots & 0 \\ F_t'^{(1)} & 0 & \dots & 0 \\ 0 & F_t'^{(2)} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & F_t'^{(m+k-1)} \end{bmatrix} a_t + H_t w_t, \tag{3.18}$$

where $F_t'^{(p)}$ denotes the first p elements of the transpose of vector F_t truncated at the first $m+k$ elements and the vector a_t is as previously defined. In equation (3.10), it was assumed each element of a_t follows a random walk. Assuming ε_t is Normally distributed with covariance S, this is a nonlinear, Gaussian state space model, and it cannot be estimated by the Kalman Filter. With the additional assumption of S block diagonal

$$S = \begin{bmatrix} \text{var}(\varepsilon_{1t}) & 0_{1\times 2} & 0_{1\times 3} & \dots \\ 0_{2\times 1} & \text{cov}\left(\left[\varepsilon_{2,t}, \varepsilon_{3,t}\right]'\right)_{2\times 2} & 0_{2\times 3} & \dots \\ 0_{3\times 1} & 0_{3\times 2} & \text{cov}\left(\left[\varepsilon_{4,t}, \varepsilon_{5,t}, \varepsilon_{6,t}\right]'\right)_{3\times 3} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

where the subscript notation i,t denotes the ith element of the vector a_t at time t, each row of the matrices in equation (3.18) is independent of the others. In addition, each row of the equation is linear. The Kalman Filter and backwards recursions are then be applied row-by-row of the equation to obtain draws for the elements of a_t.

The first row of the system corresponding is a vector of zeros, so sampling begins with the second row of the matrix equation. For the second row of the matrix equation, the state space model is

$$F_{2,t} - \Phi_{2,t} F_{t-1} = F_t'^{(1)} a_{1,t} + h_{2,t} w_{2,t}$$
$$a_{1,t} = a_{1,t-1} + \varepsilon_{1,t}.$$

The variance of $\varepsilon_{1,t}$ is the first diagonal element of *S*. For the third row of the system, the state space model is

$$F_{3,t} - \Phi_{3,t} F_{t-1} = F_t'^{(2)} \begin{bmatrix} a_{2,t} \\ a_{3,t} \end{bmatrix} + h_{3,t} w_{3,t}$$
$$\begin{bmatrix} a_{2,t} \\ a_{3,t} \end{bmatrix} = \begin{bmatrix} a_{2,t-1} \\ a_{3,t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{2,t} \\ \varepsilon_{3,t} \end{bmatrix}.$$

The relevant 2×2 block of *S* used in the Kalman Filter is the second block in the series of blocks along the diagonal of *S*, the covariance matrix $\text{cov}(\begin{bmatrix} \varepsilon_{2,t}, \varepsilon_{3,t} \end{bmatrix}')$.

For each row of the matrix equation, the Kalman Filter and backwards recursions are applied to obtain samples for the current portion of the vector a_t. Upon obtaining the sample, the corresponding block of the residual covariance *S* is updated. Following Primiceri (2005), the prior for the block of *S* related to the *i*th row of the equation is Inverse Wishart $p(S_i) = IW(i, 0.01 \times i \times I)$, where *I* is an identity matrix of dimension $i-1$. Specifying at least *i* degrees of freedom in the Inverse Wishart prior for the block of *S* with dimension $i-1$ guarantees each block of the covariance matrix is positive

definite (Hoff 2009). The conditional posterior is also Inverse Wishart

$IW\left(T-1+i,(0.01\times i\times I+U_S)^{-1}\right)$, where $U_S=\sum_t[a_t-a_{t-1}][a_t-a_{t-1}]'$.

The procedure continues until all elements of a_t have been drawn when row $m+k$ of the matrix equation in equation (3.18) is reached. Again because of the non-stationary nature of the random walk, each run of the Kalman Filter is initialized with a zero vector and diagonal covariance matrix with 0.01 on each diagonal element.

Model Comparison. In the construction of a VAR model, an adequate number of lags must be selected to capture the dynamics of the system. Additional lags can provide better fit at the cost of model parsimony and additional computational burden. In constructing the TVP-FAVAR model, a lag selection must be made for equation (3.4). The TVP-FAVAR, however, has additional complexity in that the number latent factors must also be selected. The literature applying FAVARs estimated with traditional, frequentist methods have proposed criteria such as that of Bai and Ng (2002) and the improvements to the same proposed by Alessi, Barigozzi, and Capasso (2010).

Gelman et al. (2013) discuss information criteria for Bayesian model comparison. Information criteria weigh the tradeoff between model fit and complexity. Akaike's Information Criterion (AIC) is often used in both frequentist and Bayesian applications. A Bayesian generalization of the AIC is the Deviance Information Criterion (DIC), proposed by Spiegelhalter et al. (2002, p. 584), which incorporates "…an information theoretic argument to motivate a complexity measure p_D for the effective number of parameters in a model, as the difference between the posterior mean

of the deviance and the deviance at the posterior estimates of the parameters of interest." Berg, Meyer, and Yu (2004, p. 107) present the DIC as "…easy to calculate and applicable to a wide range of statistical models." Following Berg, Meyer, and Yu (2004), the DIC can be written

$$\text{DIC} = \bar{D} + p_D, \tag{3.19}$$

where

$$\bar{D} = E\left[D(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G) \right]$$

and

$$\begin{aligned} p_D = E\left[D(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G) \right] \\ -D\left[E(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G) \right], \end{aligned}$$

where $E(\cdot)$ denotes the expectation operator, and $D(\cdot)$ is the deviance function, or

$$D(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G) = -2 \ln p(\tilde{x}, \tilde{y} \mid \tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G). \tag{3.20}$$

Substituting equation (3.20) into equation (3.19) and rearranging yields

$$\begin{aligned} \text{DIC} = -4E\left[\ln p(\tilde{x}, \tilde{y} \mid \tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G) \right] \\ + 2 \ln p\left(\tilde{x}, \tilde{y} \mid E(\tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G) \right). \end{aligned} \tag{3.21}$$

The population expectations, however, are not known and must be estimated. For the TVP-FAVAR model, the likelihood function $p(\tilde{x}, \tilde{y} \mid \tilde{f}, \Lambda, \tilde{\Phi}, \tilde{H}, \tilde{A}, R, S, Q, G)$ is complex and difficult to evaluate. Berg, Meyer, and Yu (2004) employ Kitagawa's (1996) particle filter to approximate the likelihood. The particle filter propagates a series of Z particles, or draws, representing the entire set of the model's time-varying parameters – the latent factors and the time-varying VAR coefficients, volatility, and

contemporaneous relationships – according to the dynamics specified in equations (3.3), (3.4), (3.6), (3.8), and (3.10). As they move from observation to observation, the particles are resampled with weights proportional to the likelihood of the vector of residuals from the observation equation given in equation (3.3). The complete algorithm for each iteration $b = 1,\ 2,\ \ldots,\ B$ of the MCMC samples is as follows.

1. Given the time-constant parameters of the *b*th iteration from the MCMC chain, the filter is initialized at $t = 0$ by randomly sampling values for the time-varying parameters. The starting values for the VAR coefficients, volatility, and contemporaneous relationships are each drawn from a Normal distribution centered at their corresponding time $t = 0$ mean from the *B* MCMC samples. The covariance matrices for the starting values are:
    - state vectors: a diagonal matrix of ones for the location of each factor and zeros for each known value of the observational series;
    - VAR coefficients: a diagonal matrix with 0.0001 for each diagonal element;
    - log volatility of the transition equation residuals: a diagonal matrix with 0.1 for each diagonal element; and
    - the contemporaneous relationships: a diagonal matrix with 0.1 for each diagonal element.
2. For each period $t = 1,\ 2,\ \ldots,\ T$ iterate the following steps.
    a. Stochastically propagate the VAR parameters, volatility, and contemporaneous relationships to time *t* from time $t - 1$ using the

dynamics specified in equations (3.6), (3.8), and (3.10). Once the new parameters have been propagated, calculate Ω_t using equation (3.7).

b. For each set of particles, a vector of random draws for v_t is made from the distribution of the residuals of the transition equation, that is $N(0, \Omega_t)$.

c. The predicted value of the latent state vector at time *t* is calculated for each particle using the dynamics of the transition equation described by equation (3.4) for the *t*th observation and the draws for v_t obtained in the previous step.

d. Having the predicted value of the latent state vector for each particle, the vector of residuals for the measurement equation (3.3), or

$$\begin{bmatrix} e_t \\ 0 \end{bmatrix} = \begin{bmatrix} x_t \\ y_t \end{bmatrix} - \begin{bmatrix} \Lambda_f & \Lambda_y \\ 0 & I \end{bmatrix} \begin{bmatrix} f_t \\ y_t \end{bmatrix},$$

is calculated from the relationship described in equation. For the *z*th particle, the joint density of the residuals, $N(0, R)$, is evaluated at the vector of predicted residuals to obtain the weights ω_{zt}, where $z = 1, 2, \ldots, Z$.

e. The particles are resampled with weights ω_{zt}, and the algorithm continues to the next *t* by returning to step (a) and continuing the filter with the resampled particles.

Upon completion of the particle filtering algorithm, the log-likelihood for the *b*th MCMC iteration is estimated

$$(3.22) \qquad \ln p(\tilde{x}, \tilde{y} \mid \tilde{f}_b, \Lambda_b, \tilde{\Phi}_b, \tilde{H}_b, \tilde{A}_b, R_b, S_b, Q_b, G_b) \approx \sum_{t=1}^{T} \left[\ln \left(\frac{\sum_{i=1}^{Z} \omega_{it}}{Z} \right) \right] \equiv \hat{l}_b.$$

The algorithm is repeated for each iteration *b* of the MCMC procedure. By the Law of Large Numbers, the DIC can then be estimated using the sample analog of the population expectation in equation (3.21) so that

$$(3.23) \qquad \text{DIC} \approx -4 \frac{\sum_{b=1}^{B} \hat{l}_b}{B} + 2 \ln p\left(\tilde{x}, \tilde{y} \mid \bar{\tilde{f}}, \bar{\Lambda}, \bar{\tilde{\Phi}}, \bar{\tilde{H}}, \bar{\tilde{A}}, \bar{R}, \bar{S}, \bar{Q}, \bar{G} \right),$$

where the second term is the likelihood of the measurement equation (3.3) evaluated at the sample mean of the parameters (denoted by the bar notation) across all *B* draws of the MCMC. This, too, can be estimated via particle filter. Kitagawa (1996) suggests $Z = 1,000$ is sufficient for estimating the mean of the distribution.

Postestimation Analysis. The informational series contain a broad variety of series describing the upstream, midstream, and downstream of the petroleum industry and the US economy. The TVP-FAVAR setup allows each of these series to be examined individually to understand its dynamics and their evolution. As mentioned previously, one strength of the FAVAR framework is the ability to include many time series while still being able to use the traditional tools of VAR analysis, such as IRFs and FEVDs.

In traditional VAR analysis, the IRFs are commonly calculating using the moving-average (MA) representation of the system (Lutkepohl 2005). The MA representation transforms the system such that the VAR in equation (3.5) is a function of lags of the innovation noise

(3.24) $$F_t = v_t + \Theta_{1,t} v_{t-1} + \Theta_{2,t} v_{t-2} + \ldots = \Theta_t(L) v_t,$$

where $\Theta_t(L)$ is a lag polynomial operator and, like the original VAR coefficient matrices Φ_t from equation (3.5), is time-varying. A *p*-step-ahead IRF, as defined in Hamilton (1994), is the vector of changes in the values of the variables in the VAR system given a unit shock in one of the innovations in v_t, or

(3.25) $$\frac{\partial F_{t+p}}{\partial v_t},$$

where *p* is a nonnegative integer denoting the time steps ahead. The IRF shows how a shock in one of the VAR's variables at a given period affects the values of the variables for future time periods. Using the MA representation from equation (3.24), the IRF can be shown to be

(3.26) $$\frac{\partial F_{t+p}}{\partial v_t} = \frac{\partial\left[v_{t+p} + \Theta_{1,t} v_{t+p-1} + \Theta_{2,t} v_{t+p-2} + \ldots + \Theta_{p,t} v_t + \ldots \right]}{\partial v_t} = \Theta_{p,t},$$

where the item at the *i*th row and *j*th column is the effect of a shock in the *j*th variable on the *i*th variable in the vector F_{t+p}.

In economic data, the innovations in v_t, however, are not necessarily independent, so a shock in one variable could happen concurrently with a shock in another variable. Recall that v_t is Normally distributed with covariance matrix Ω_t. The IRFs as currently defined "…may provide a misleading picture of the actual dynamic relationships between the variables" (Lutkepohl 2005, p. 57). The common practice is to

calculate structural IRFs (SIRFs) in which the innovations to each variable are independent of each other.

Gambetti (2012) provides the orthogonalized version of equation (3.24) under TVP as

$$\begin{aligned} F_t &= \Theta_t(L) C_{\Omega_t} W_t v_t \\ &= C_{\Omega_t} W_t v_t + \Theta_{1,t} C_{\Omega_{t-1}} W_{t-1} v_{t-1} + \Theta_{2,t} C_{\Omega_{t-2}} W_{t-2} v_{t-2} + \ldots , \end{aligned} \tag{3.27}$$

where C_{Ω_t} is the Choleski decomposition of the time-varying residual covariance matrix Ω_t and W_t is the rotation matrix determining the ordering of the series. The SIRF at p steps ahead of time t can be obtained using the orthogonalized MA representation in equation (3.27) to find $F_{t+p} = \Theta_{t+p}(L) C_{\Omega_{t+p}} W_{t+p} v_{t+p}.$ and then differentiating to obtain

$$\frac{\partial F_{t+p}}{\partial v_t} = \Theta_{p,t+p} C_{\Omega_t} W_t. \tag{3.28}$$

The pth matrix of $\Theta_t(L)$ is the only matrix in the MA representation needed to calculate the SIRF, and it can be obtained by substitution (Gambetti 2012). For $p = 1,$ the one-step-ahead forecast is $F_{t+1} = \Phi_{t+1} F_t + v_{t+1}.$ By substituting $\Phi_t F_{t-1} + v_t$ for F_t the expression becomes

$$\begin{aligned} \hat{F}_{t+1} &= \Phi_{t+1} F_t + v_{t+1} \\ &= \Phi_{t+1} \left[\Phi_t F_{t-1} + v_t \right] + v_{t+1} \\ &= \Phi_{t+1} \Phi_t F_{t-1} + \Phi_{t+1} v_t + v_{t+1}. \end{aligned}$$

The needed coefficient matrix is $\Phi_{t+1}.$ Similarly, using the two-step-ahead forecast to obtain the matrix for $p = 2$ provides

$$\begin{aligned}\hat{F}_{t+2} &= \Phi_{t+2}\hat{F}_{t+1} + v_{t+2} \\ &= \Phi_{t+2}\left[\Phi_{t+1}F_t + v_{t+1}\right] + v_{t+2} \\ &= \Phi_{t+2}\left[\Phi_{t+1}\left(\Phi_t F_{t-1} + v_t\right) + v_{t+1}\right] + v_{t+2} \\ &= \Phi_{t+2}\Phi_{t+1}\Phi_t F_{t-1} + \Phi_{t+2}\Phi_{t+1}v_t + \Phi_{t+2}v_{t+1} + v_{t+2}.\end{aligned}$$

In this case, the coefficient matrix is $\Phi_{t+2}\Phi_{t+1}$. Following this pattern, for a *p*-step ahead SIRF, the coefficient matrix from the MA representation is $\Phi_{t+p}\Phi_{t+p-1}\ldots\Phi_{t+2}\Phi_{t+1}$.

Upon calculating the SIRFs for the vector F_t, the SIRFs for informational series can be obtained. Given $x_{t+p} = \begin{bmatrix}\Lambda_f & \Lambda_y\end{bmatrix}F_{t+p} + e_{t+p}$, the time-varying SIRFs for x_{t+p} are

$$\begin{aligned}\frac{\partial x_{t+p}}{\partial v_t} &= \begin{bmatrix}\Lambda_f & \Lambda_y\end{bmatrix}\frac{\partial F_{t+p}}{\partial v_t} \\ &= \begin{bmatrix}\Lambda_f & \Lambda_y\end{bmatrix}\Theta_{p,t+p}C_{\Omega_t}W_t.\end{aligned} \tag{3.29}$$

Many techniques to identify the rotation matrix W_t exist, such as recursive ordering schemes and graph-based approaches. An "agnostic" approach proposed by Uhlig (2005) is the placement of restrictions on the signs of the SIRFs. In this method, potential SIRFs are calculated and then evaluated compared to theoretical responses. For example, Uhlig (2005) identifies a monetary policy shock such that the SIRFs of the price level and nonborrowed reserves are not positive while the SIRFs for the federal funds rate are not negative for six months. Lippi and Nobili (2008) employ a sign restriction identification scheme to identify shocks in oil supply and demand as well as aggregate demand and supply. Amir-Ahmadi and Uhlig (2009) employ sign restrictions to identify a monetary policy shock as an extension of Bernanke, Boivin, and Eliasz's (2005) original FAVAR. Juvenal and Petrella (2012) investigate speculation for oil

using an FAVAR with SIRFs identified by sign restrictions. The TVP-FAVAR of Mumtaz, Zabcyk, and Ellis (2011) is also identified by sign restrictions.

The literature using sign identification for VARs while examining crude oil commonly identify three shocks that are applied here (Kilian and Murphy 2010, Juvenal and Petrella 2012, Melolinna 2012, Baumeister and Peersman 2013). The first shock is a negative oil supply shock, "…defined as any unanticipated shift in the oil supply curve that results in an opposite movement of oil production and the real price of crude oil. During an oil supply disruption inventories are depleted in an effort to smooth oil production and real activity contracts" (Juvenal and Petrella 2012, p. 13). In the context of the TVP-FAVAR, for this shock, the following restrictions are applied:

- real oil prices increase contemporaneously with the shock;
- field production of crude oil and crude oil stocks at tank farms decrease contemporaneously with the shock;
- the consumer price index (CPI) and producer price index (PPI) increase with the shock; and
- real activity, measured by the industrial production index (IP), decreases contemporaneously with the shock.

The second shock is an aggregate demand shock (Juvenal and Petrella 2012), which increases the demand across the economy, raising the price level. Because of the increased demand in the economy, oil demand also increases, causing oil prices and production to increase. In the context of this TVP-FAVAR, this shock necessitates:

- IP increases contemporaneously with the shock;

- CPI and PPI increase with the shock; and
- real oil prices and field production of crude oil increase contemporaneously with the shock.

The third shock is an oil demand shock in which:

- real oil prices and field production of crude oil increase contemporaneously with the shock;
- IP decreases contemporaneously with the shock; and
- CPI and PPI increase with the shock.

Kilian and Murphy (2010) study the petroleum industry and argue sign restrictions alone are insufficient for identification of a VAR. In a purely sign-identified model, the price elasticities of oil suppliers in the same month as the shock can be large in magnitude, while the literature is largely in agreement that the short-run elasticity of oil supply is close to zero. They propose an additional restriction limiting the elasticity of oil production to 0.0258 in the period of the shock for the two demand shocks. In addition, the restrictions are only required to be valid at the contemporaneous period and not at any future horizon. This additional restriction is used here.

Sign identification is performed following the algorithm of Rubio-Ramirez, Waggoner, and Zha (2010) and as outlined in Kilian and Lutkepohl (2017) and Baumeister and Peersman (2013). For each period $t = 1, 2, \ldots, T$, the following are iterated.

1. For the current period *t*, draw at random one of the saved iterations from the MCMC output. Given the values of the time-varying parameters from this draw

at *t*, the time-varying VAR coefficients, residual volatility, and contemporaneous relationships are propagated forward in time using the transition dynamics described in equations (3.6), (3.8), and (3.10). In this way, all potential sources of uncertainty determining the SIRFs are considered (Baumeister and Peersman 2013).

a. Given the simulated time-varying parameters, draw an $(m+k)\times(m+k)$ proposal matrix of independently distributed standard Normal variables. Use the *QR* decomposition with necessary modifications (Kilian and Lutkepohl 2017) to obtain an orthonormal matrix *O* from the proposal matrix of random deviates. By drawing the matrices in this manner, the proposal impact matrix *O* is uniformly distributed over the space of orthonormal matrices.

b. The proposed zero-step SIRF for the VAR equation is $C_{\Omega_t}O'$. The Cholesky decomposition here is not used for identification but rather to orthogonalize the residuals, so the ordering of the variables is not important. Calculate $\begin{bmatrix}\Lambda_f & \Lambda_y\end{bmatrix}C_{\Omega_t}O'$ to obtain the SIRFs for the informational series. For each of the restricted responses series, verify the sign and impact elasticity of the responses in $\begin{bmatrix}\Lambda_f & \Lambda_y\end{bmatrix}C_{\Omega_t}O'$ agree with the restrictions. None of the inequalities tested are strict. If the restrictions are satisfied, the draw is saved, and the SIRFs are calculated.

If they are not satisfied, the draw is discarded, and the algorithm begins again.

2. The proposal-rejection step is repeated until 100 realizations of the impact matrix have been accepted for each set of coefficients drawn in the first step at each *t*.
3. The mean SIRF across the 100 accepted realizations is calculated and saved for this set of parameters.

The entire preceding procedure is performed 100 times for each *t*, using a random draw for the parameters from the MCMC output for each realization.

In VARs estimated using Bayesian techniques and identified via sign restrictions, the highest posterior density (HPD) interval and median for each SIRF are commonly reported to describe the distributions for potential responses. The median and HPD, however, only represent features of the set of possible SIRFs and not actual realizations of draws of the dynamics of the variables following a shock (Killian and Murphy 2010, Uhlig 2017). In addition, the HPD is constructed assuming each realized SIRF is equally likely, but Killian and Murphy (2010) show this is not necessarily true. In their application, they successfully reduce via sign and elasticity restrictions the number of accepted IRFs to a smaller number that are qualitatively similar. In this dissertation's model, however, the same kind of reduction is not achieved; sign and elasticity restrictions identified a set of IRFs of which the individual IRF realizations potentially suggest different conclusions about the magnitude and shape of the responses to shocks. Because of the difficulty in selecting a representative model and to simplify comparison across the literature, the median and HPD intervals are reported.

Gambetti (2012) also provides formulas for the time-varying FEVDs. Let the *p*-step SIRF at time *t* of the *i*th informational series with respect to the *j*th shock be the entry of at row *i*, column *j* of the matrix

(3.30) $$\left[\left[\Lambda_f \quad \Lambda_y\right]\Theta_{p,t+p}C_{\Omega_t}W_t\right]_{i,j},$$

and denote this value $\theta^{i,j}_{p,t+p}$. The *p*-step time-varying FEVD for informational series *i* with respect to shock *j* at time *t* is then

(3.31) $$\frac{\sum_{d=0}^{p}\left(\theta^{i,j}_{d,t+d}\right)^2}{\sum_{i=1}^{m+k}\sum_{d=0}^{p}\left(\theta^{i,j}_{d,t+d}\right)^2}.$$

Here, the FEVDs for each time period are calculated using the output of the SIRF identification procedure.

Simulation Details. Models specified with different combinations of the number of latent factors and the number of lags included in the VAR are estimated. The limiting factor for the MCMC estimation and analysis of the TVP-FAVAR model is a system's memory. Saving the results for convergence analysis and traditional VAR analysis requires storing many iterations of hundreds of parameters. Simulations are written in R (R Core Team 2018) and performed using high-performance computing resources at Texas A&M University.

While the literature provides clear suggestions for prior specification, substantial experimentation was performed using less informative priors (such as Jeffrey's) than those used here. For most of the prior specifications attempted, there were many issues with numerical stability in the Kalman Filter caused by the inability to invert the

covariance matrix. The two most important prior specifications for attaining results appear to be *Q*, the covariance matrix governing the evolution of the time-varying VAR parameters, and the prior covariance for the observation equation coefficients. Weak priors on both led to explosive behavior in the estimated parameters. Explosive behavior in the time-varying VAR coefficients, residual volatility, or contemporaneous relationships also led to degeneracy in the positive definiteness of covariance matrices, making numerical evaluation of matrix inverses impossible. Primiceri (2005) also notes the importance of prior choice, especially in the case of the time-varying coefficients.

As such, other strategies were adopted to improve numerical stability of the estimation algorithm. First, the Joseph form of the Kalman Filter prediction covariance update ensures a numerically symmetric covariance matrix is obtained for each observation of the state vectors (McClure 2016). The Joseph update replaces the final step of each iteration in equation (3.12) with $P_{t|t} = (I - K_t \begin{bmatrix} \Lambda_f & \Lambda_y \\ 0 & I \end{bmatrix}) P^F_{t|t-1} (I - K_t \begin{bmatrix} \Lambda_f & \Lambda_y \\ 0 & I \end{bmatrix})' - K_t R K_t'$. Matrix inversion was avoided as it is generally not computationally efficient nor precise (Cook 2010).

To analyze convergence, each run of the simulation is performed in parallel with two chains of 60,000 iterations, discarding the first 30,000. The starting values of the two chains are:

1. all coefficients are set to zero, all variance parameters are set to one, and all covariance parameters are set to zero; and

2. all values are drawn randomly from an "overdispersed" uniform distribution. The diagonal elements of *Q*, *G*, and *S* are drawn from the range [0.0001, 10]; the diagonal elements of *R* are drawn from [0.0001, 10]; the time-varying VAR volatility are drawn from [0.0001, 4]; the VAR coefficients are drawn from [-2, 2]; the lower-triangular elements of *A* are drawn from [-5, 5], and the coefficients in the observation equation are set to zero.

Because of the large number of parameters – especially those that are time-varying – in the model, the output of long chains cannot be feasibly saved for analysis. In models with many parameters, Gelman et al. (2013) recommend thinning. Here, every 15^{th} iteration is kept, yielding 2,000 draws from the posterior. The use of thinned chains also reduces autocorrelation of the samples in the smaller portion that is kept for analysis. Assessing convergence is difficult due to the large number of parameters. The chains are monitored for convergence using a combination of visual inspection, the effective sample size, and the scale reduction factor (Gelman et al. 2013).

Data

Monthly data used for estimation include two broad categories of time series: industry-specific measures for the petroleum industry in the US and macroeconomic variables for the overall economy. Data cover the period January 1993 through March 2018. In the first set of informational series, variables pertinent to each sector of the petroleum industry are included. As indicated by the literature, rig counts and field production are commonly used to gauge upstream activity. For the downstream, measures of refinery production and utilization are also included. Volumes of storage, imports and exports

are barometers of midstream activity. The petroleum industry data are obtained from the United States Energy Information Administration (2018).

The second type of informational series includes macroeconomic variables comparable to those included in Zagaglia (2010), Amir-Ahmadi and Uhlig (2009), An, Jin, and Ren (2014), Ipatova (2014) and Binder, Pourahmadi, and Mjelde (2018). These series include measures of money supply, inflation, interest rates, employment, wages, production measures, exchange rates and international trade flows, housing activity, and price deflators. These data are obtained from the Federal Reserve Bank of St. Louis (2018).

A full list of the series employed by the model is included in Table 3.1. All dollar amounts are deflated using the CPI. For each series, the table provides the source of the data, and the transformation used, if any, either to dampen visually observed heteroskedasticity (log transformation) or to induce stationarity (first difference or first difference of log transformation). These are the same transformations used by Bernanke, Boivin, and Eliasz (2005). Upon application of a transformation (if applicable), each series is standardized by subtracting the mean and dividing by its standard deviation.

As mentioned previously, the two observational series are oil prices (first difference of log transformation) and the federal funds rate. A plot of oil prices is shown in Figure 3.1. Shaded areas represent recessions (National Bureau of Economic Research 2018). In the sample, there are several interesting time periods, including 1) June 2008, the all-time high oil price; 2) February 2009: the lowest oil price during the

Great Recession, a 70 percent decrease in price in less than one year; and 3) February 2016: most recent low price. A plot of federal funds rate is shown in Figure 3.2. Historically low values of the federal funds rate, used to stimulate the economy during the Great Recession and present through 2017, are easily spotted relative to the higher rates of the 1990s and mid-2000s.

Results

The DIC value for each estimated specification of the TVP-FAVAR is shown in Table 3.2. Combinations of lags and factors were selected for computational feasibility. Some of the DIC values for the three- and four-factor specifications are large compared to the others. These DIC values may reflect issues with convergence of the MCMC chains, as examination of the chains for the three-factor, two-lag specification suggested longer chains would be beneficial.

One latent factor is preferred by the DIC. The model with one latent factor and three lags in the VAR is suggested as the best tradeoff between fit and complexity by the DIC. The DIC values for the models with one latent factor and either one or 12 lags are similar to the one-factor, three-lag specification. Hamilton and Herrera (2004) argue the effects of oil price shocks are understated in VAR models with short lag lengths. In their survey of the literature, they find, for studies using quarterly data, the two lags with the largest coefficients are usually lags three and four. They suggest longer lag lengths are necessary to control for seasonality and argue Bernanke, Gertler, and Watson (1997) do not include sufficient lags. While 90-percent HPDs for the time-varying lag-12 autoregressive coefficient of the latent factor contains zero, 80-percent HPDs for these

coefficients do not. The equivalent 80- and 90-percent HPDs for the 12^{th} lag autoregressive coefficients from the equations for oil price and interest rate, however, contain zero. Given the similar DIC values and the potential importance of lag 12 for the latent factor, the model with one latent factor and 12 lags in the TVP-VAR is selected for analysis. Results and inference only from this model are presented.

For the selected model and chain, parameters with relatively small effective sample sizes include the latent factor and the time-varying matrix H_t (at some, but not all, periods), and parts of the diagonal covariance matrix G. Because two chains were run simultaneously, convergence can be assessed with the potential scale reduction factor (Gelman et al. 2013). This metric decreases to one as the number of iterations in the MCMC goes to infinity. Of the parameters evaluated, the diagonal elements of R in general have the smallest values of the potential scale reduction factor, and those for the coefficients of the observation equation are also relatively small. In general, the highest values found during evaluation are those for the elements of G, with the largest value at 1.643.

The Latent Factor

A plot of median of the latent factor's evolution over time is shown in Figure 3.3 along with 68- and 90-percent HPDs. The practice of showing 68- and 90-percent HPDs is followed from the Bayesian VAR literature (Primiceri 2005, Mumtaz, Zabczyk, and Ellis 2011, Amir-Ahmadi and Uhlig 2009, Lippi and Nobili 2008). For the latent factor, both HPDs are almost indistinguishable from the median. A few distinct periods of economic history are suggested by the plot. The first period runs from the start of the

data in the mid-1990s through 2000. In this period, the latent factor is relatively larger than in the early 2000s and after the Great Recession. Around the 2001 recession, the factor decreases, but following the recession, it trended upward until about 2005. At the beginning of the Great Recession, the value of the latent factor drops but returns to its 2008 levels values by 2010. Since 2010, the latent factor is relatively constant. The latent factor also demonstrates seasonality. This partially explains why the lag-12 autoregressive coefficient for the latent factor in the TVP-VAR does not contain zero in the HPDs compared to the autoregressive coefficients of the oil price and interest rate equations. This is expected as most of the informational series are not seasonally adjusted.

An important consideration in evaluating the factor model is the fit provided by the factors. A simple method to determine which series are loaded onto the factor is the R-squared value from an OLS regression of each informational series individually on the latent factor. A graphical summary of the series-level fit is provided in Figure 3.4. For each informational series, the median R-squared value across the kept posterior draws is plotted. The series are grouped into broad categories to study patterns in the latent factor for different types of variables (Table 3.3). The R-squared values for series categorized as petroleum-related prices, which include spot prices for gasoline and futures prices for crude oil, are very small, as are R-squared values for series categorized as related to the federal funds rate. The Bayesian estimation of the model provides joint estimation of the latent factor and its loading. The low R-squared values for regressions involving informational series closely related to oil prices and the federal funds rate likely occurs

because the full set of regressors for each informational series in the TVP-FAVAR model includes both oil prices and the federal funds rate in addition to the latent factor. Because these series are explained well by oil prices, the "uncovering" of the latent factor appears to place more weight on the categories series less related to the oil prices and the federal funds rate. Many of the R-squared values obtained using industrial production series as regressor are relatively larger. This is likely due to the restrictions necessary for identification of the factors which placed IP first in the ordering of all series. Across the various industry-specific IP series, however, there is a substantial amount of variation in R-squared values, suggesting the latent factor is not closely tied to all of these.

Time-Varying Residual Volatility

The time-varying VAR median residual volatility is plotted for each equation in the VAR equation, latent factor, oil price, and federal funds rate, in Figure 3.5 along with 68- and 90-percent HPDs. While visual inspection of the plots suggests time-varying volatility for oil prices and the federal funds rate, the residual volatility of the latent factor is relatively constant, potentially decreasing slowly over time. Oil price volatility generally increases from 1994 through the early 2000s where the volatility becomes relatively constant. An increase in oil price volatility occurs between 2008 and 2009, corresponding to the 70 percent decrease in prices over the same period. Oil price volatility again increases in 2016 as oil prices fell again. The volatility in the federal funds rate increases in 1994, 2001, and 2008. In each of these periods, the federal funds rate either increased or decreased over a short period of time.

Dynamic Effects of the Structural Shocks

Several series are selected for analysis and discussion to understand the dynamics of the petroleum industry over the sample period. For each selected series, three plots are provided to analyze the effects of the shocks over time:

1) a plot showing the cumulative effect of the shocks at one year after the shock against time along with 68- and 90-percent HPDs;
2) a plot showing the FEVD at the 12-month horizon against time along with 68- and 90-percent HPDs; and
3) a plot showing the time to stabilization (TTS) in months following the shock.

For each period, the plots show values obtained from the time-varying SIRFs and FEVDs using that period's values of the time-varying parameters as the base for the propagation procedure described previously. To calculate TTS of a shock, a 90-percent HPD is constructed for the instantaneous SIRF at each horizon from zero to 42 steps ahead. The first horizon with an HPD interval containing zero is considered the period of stabilization. Each of the duration plots contains an overlay of a 12-month arithmetic moving average.

These plots provide insights into the evolution of the dynamics of the industry and the economy. In the text that follows, the median values are discussed unless otherwise noted. In addition, the median cumulative change in each series in response to the shocks at the period contemporaneous to the shock and one year ahead are shown for selected series in Table 3.4. Appendix A contains plots of instantaneous IRFs for these series at selected periods. No results are presented on the federal funds rate; the federal

funds rate was included as an observational series for its importance in policy, but it is not analyzed.

Oil Prices. By the sign restrictions imposed for identification, oil prices must not decrease concurrently with a shock in oil supply, aggregate demand, or oil demand. A plot of the time-varying 12-month cumulative response of oil prices to the shocks is displayed in Figure 3.6. The effect of oil supply shocks is relatively constant over time but is marked by increases in 2001, 2008, and 2009. The evolution of the effects of aggregate demand shocks (Figure 3.6) is similar in shape to that of the residual volatility for oil prices (Figure 3.5). The magnitude of responses to shocks in aggregate demand shows a generally upward trend between 1994 and 2008. To compare, the median cumulative increase in oil prices one year after aggregate demand shocks in December 1994 and December 2008 are 20.9 and 37.9 percent. Following the Great Recession, the responses of oil prices decreased in magnitude to their mid-1990s levels. The cumulative effects of oil demand shocks demonstrate the most variability of the three 12-month ahead cumulative responses. The effects of oil demand shocks generally evolve opposite of the oil supply shocks; as the responses of oil prices to oil supply shocks increases, the responses to oil demand shocks decreases.

The relative importance of shocks over most of the sample period, measured by the magnitude of the change in oil prices, is oil demand, aggregate demand, and oil supply shocks. This ordering is consistent with Kilian and Murphy (2010). Both the 2001 and 2008 recessions are preceded approximately half a year by rapidly decreasing responsiveness of oil prices to shocks in oil demand. Shortly after the beginning of both

recessions, there is a jump in price responsiveness to oil supply shocks. There is a reduction in the magnitude of oil price responses to shocks in aggregate demand and oil demand beginning in 2016, but there have been no changes in the magnitude of responses to supply shocks.

A plot of the time-varying proportion of the 12-month forecast error variance for oil prices is shown in Figure 3.7. Oil demand shocks also account for the largest proportion of forecast error variance in oil prices. For the majority of the sample, oil demand shocks explain at least half of the forecast error variance for oil prices, followed by aggregate demand shocks and oil supply shocks. During the 2001 and 2008 recessions, the proportion of variance explained by oil demand shocks falls, while the proportion explained by oil supply shocks increases. This change is similar to the change exhibited in the 12-month cumulative responses shown in Figure 3.6.

TTS of oil prices after the shocks has been relatively constant (Figure 3.8). Price changes continue about one year following an oil supply shock. In the early 1990s, TTS is generally 12 months, but since 2006, shocks often resolve in less than 12 months. TTS following an aggregate demand shock is over one year. TTS for oil demand is similar to TTS for aggregate demand shocks, but at the beginning of the Great Recession, TTS dropped temporarily.

Oil Production (Upstream). By the sign restrictions imposed for identification, oil production also responds concurrently to all three shocks. Unlike those of oil prices, the responses of oil production to the shocks generally exhibit small variations over time (Figure 3.9). For an oil supply shock, the plot shows a brief increase in the magnitude of

the responses during the 2001 and 2008 recessions. Note the shock is negative, so an increase in magnitude of the shock implies the response is more negative. After decreasing in magnitude at the beginning of the Great Recession, the effects of aggregate demand shocks increase in magnitude between 2008 and 2009. Responses to oil demand shocks increase slightly in magnitude during the recessions.

Kilian and Murphy (2010 p. 12), following Hamilton (2009), argue "Oil producers will respond to unanticipated oil price increases only if that increase is expected to persist." The difference in oil producers' responses to aggregate demand shocks and oil demand shocks suggests producers perceive unanticipated aggregate demand shocks as relatively more permanent than unanticipated oil demand shocks. For most of the sample, aggregate demand shocks explain over half of the forecast error variance (Figure 3.10). The proportion of forecast variance explained by oil supply shocks increases abruptly in 2001, 2008 and 2009 in a similar manner to the increases seen in the proportions of forecast error variance for oil prices.

Over time, TTS of oil supply shocks decreases from around one year in the mid-1990s to under two quarters in 2018 (Figure 3.11). At the beginning of the recessions, TTS of oil supply shocks increases by about one quarter. TTS of aggregate demand shocks increases over time. Except during recessions and prior to 2008, aggregate demand shocks appear to resolve in approximately two quarters, but after 2009, TTS increases, and shocks in aggregate demand stabilize at approximately three quarters. Oil demand shocks generally resolve in one quarter, but during the 2008 recession, TTS increases to almost two quarters.

Oil Storage (Midstream). Time-varying cumulative responses of oil storage 12 months following the shock are presented in Figure 3.12. The effects of oil and aggregate demand shocks are relatively larger in magnitude compared to oil supply shocks. While the 12-month cumulative effect of an oil supply shock on storage is relatively constant over the sample period, the magnitude of the responses increases early in the two recessions. The effects of the two demand shocks again evolve like the volatility of oil prices, increasing in magnitude over the 1990s and decreasing in magnitude following the Great Recession. Unlike the responses to shocks in oil supply and oil demand, the responses to aggregate demand shocks demonstrate no jumps during either recession. Storage is less responsive to oil demand shocks during the recessions.

In the following discussion, consider decreasing oil prices (Figure 3.1) and the responses of oil production (Figure 3.9) and oil storage (Figure 3.12) during the two recessions. During recessions, shocks in oil supply generate larger decreases in crude oil production than in normal periods. This is driven by both reduced economic activity as well as a decrease in the profitability of oil production due to decreasing oil prices. Because the reduction in crude oil production is relatively larger, stores of crude oil are depleted more quickly in recessions than in other periods.

Examination of the time-varying proportions of 12-month forecast error variance explained by the shocks (Figure 3.13) reveals an increase in the proportion explained by oil supply shocks during the recessions, especially at the beginning of the Great Recession. In both recessions, the explanatory role of oil demand shocks for oil storage diminished. TTS following an oil supply shock is generally six months (Figure 3.14).

Stabilization following aggregate demand shocks is not achieved until over 12 months. TTS for oil demand shocks is under 12 months but decreases to about six months during the recessions.

Refinery Capacity Utilization (Downstream). A plot of 12-month cumulative changes in refinery capacity utilization is provided in Figure 3.15; the FEVD for 12 months ahead is shown in Figure 3.16. The reduction in utilization after an oil demand shock follows from the increased acquisition costs for oil, leading to higher prices for refined goods, causing the quantity demanded of their products to decline. While the 12-month ahead change in refinery capacity utilization shows little change over time, the FEVD shows some change, as oil demand shocks account for a larger share of forecast error variance in the mid-1990s and during the two recessions. Like oil storage, refinery capacity utilization demonstrates a reduction in responsiveness to oil supply shocks and an increase in responsiveness to oil demand shocks during recessions. TTS following an oil supply shock is less than 12 months but decreases during the recessions (Figure 3.17). Through the early 2000s, the effects of aggregate demand shocks are usually stabilized in less than 12 months, but since the mid-2000s, TTS is six months or less. Oil demand shocks stabilize within six months or less. This may be because oil demand shocks are likely the result of precautionary measures of producers and are not driven by changes in demand for consumption of refiners' final products.

Gasoline Sales by Refineries (Downstream). The dynamics of refiners' sales of gasoline to retail outlets are relatively time invariant; neither the plots of 12-step ahead cumulative changes (Figure 3.18) nor the 12-step ahead FEVD (Figure 3.19) show

changes over time. For all shocks, TTS is less than one year (Figure 3.20) and has been generally decreasing since the mid-1990s. During recessions, TTS decreases for oil supply shocks.

Industrial Production (Broader US Economy). For IP, the cumulative response after one year is relatively constant (Figure 3.21), and the 12-step ahead FEVD shows little variation (Figure 3.22). During recessions, the effects of oil supply shocks decrease in magnitude slightly, while those of oil demand shocks increase slightly in magnitude. Of the three shocks, aggregate demand shocks and oil supply shocks are associated with the largest responses in magnitude after one year. Most of the HPDs for cumulative responses to oil demand shocks after one year contain zero, implying IP reverts to its pre-shock levels within a year. Oil demand shocks to IP are generally stabilized at one quarter but longer during recessions (Figure 3.23). TTS for aggregate demand shocks decreases since the 1990s.

Producer Price Index (Broader US Economy). In Figure 3.24, a plot of the cumulative responses of PPI after one year is provided. PPI increases in response to the three shocks. Oil demand shocks generate the largest increases in PPI, followed by aggregate demand and oil supply shocks. This is likely because oil prices respond in magnitude most to oil demand shocks, so the increase in oil prices following an oil demand shock likely causes PPI to increase. The effects of the shocks, especially the aggregate demand and oil demand shocks, display variation over time. An oil demand shock in December 1994, for example, would yield a median increase in PPI of 1.9 percent after one year, while a shock in December 2000 would generate a 3.1 percent

increase. For the two demand shocks, the 12-month responses in PPI increase throughout the 1990s and until the 2001 recession. For oil demand shocks, the magnitude of the responses decreases prior to the recessions. A small increase in the magnitude of responses to oil supply shocks is also characteristic of the beginnings of the recessions. The 12-month FEVD (Figure 3.25) also demonstrates this pattern. TTS of the shocks are shown in Figure 3.26. Aggregate demand shocks, which produce the responses that stabilize slowest, demonstrate stabilization at five quarters, followed by oil demand shocks (four quarters) and oil supply shocks (two quarters). There is more variability in TTS following an oil supply shock compared to the other two shocks.

Discussion and Conclusions

A TVP-FAVAR model (Mumtaz, Zabczyk, and Ellis 2011) is applied to both capture information contained in many data series and allow for changes in the dynamic responses of the variables in the system. This model is used to study components of the upstream, midstream, and downstream sectors of the petroleum industry by examining the effects of hypothesized shocks in oil supply, aggregate demand, and oil demand as identified through sign restrictions. There is visual evidence of time-varying residual volatility for oil prices and the federal funds rate, but the residual volatility for the latent factor is relatively constant. For oil prices specifically, the findings here are similar to those of Kilian and Murphy (2010); oil prices are most responsive in terms of magnitude to shocks in oil demand, followed by aggregate demand and oil supply. On the other hand, for crude oil production, shocks in aggregate demand commonly have a larger impact than oil demand shocks. For series displaying time variation in the cumulative

responses, the evolution of the magnitude of these effects is similar to the evolution of oil price volatility. The effects of shocks in oil supply increase in magnitude during recessions.

The results suggest oil prices and oil storage at tank farms (midstream) have evolved in their responses over the last 25 years, but the dynamics of field production (upstream), refinery capacity utilization (downstream), and gasoline sales (downstream) show little visual evidence of evolution in the profiles of their responses. While the magnitude of the responses of field production of crude oil, refinery capacity utilization, and gasoline sales are relatively constant as measured by magnitude, TTS shows some variability, many of the responses are affected by recessions. For the economy as a whole, the Industrial Production Index also appears relatively constant in its dynamics over the sample, but the dynamics of the Producer Price Index are exacerbated by increased oil price volatility.

As a way of addressing the use of large datasets in VAR analysis, the TVP-FAVAR model incorporates many variables while providing methods for innovation accounting with time variation. It now may be feasible to include hundreds of series in the VAR framework, but with new methodologies, there are new, or at least different, costs. One cost of the Bayesian TVP-FAVAR is the difficulty in prior selection given its importance in obtaining numerically stable results. A second is the tradeoff between model parsimony and fit. Third is the practical application of the sign identification procedure. Kilian and Murphy (2010) reduce the number of acceptable VAR models to a small set displaying similar dynamics using sign and elasticity restrictions. In this

application, a much broader set of SIRFS was accepted during identification, motivating the reporting of medians and HPDs. This makes interpretation of the dynamics of any one realization impossible. The set of accepted models could potentially be further limited by imposing additional sign or elasticity restrictions, but identifying an FAVAR via sign restrictions is challenging. Although the researcher can impose theoretical restrictions on the dynamics of any of the informational series, the placement of restrictions in the FAVAR context is difficult because of the vast number of series (potentially hundreds). With additional restrictions, it can become more difficult to accept rotation matrices. And while it must be used to guide the choice of restrictions, economic theory does not distinguish between different but closely related series. Examples of these are the CPI and PPI series. There are many variations of these series disaggregated by type of good or industry. A shock in aggregate demand increases prices in the economy, but which series should be restricted? This is generally not an issue encountered in the sign identification of small-scale VARs. Uhlig's (2017, p. 99) first two principles guiding the placement of sign restrictions are "If you know it, impose it!" and "If you do not know it, do not impose it!"

Further research should continue to investigate model performance and prior selection. Because the residual volatility of the latent factor is relatively constant, it may be beneficial to restrict this parameter to be a known constant to simplify the model. An additional means of simplifying the parameterization of the TVP-FAVAR is to consider the observation equation of the state space model as a seemingly unrelated regressions-type model with each equation's regressors being determined by Bayesian methods. In

attempting this formulation, it became apparent the computing time required to estimate such a model rendered it infeasible: another realization of the curse of dimensionality. For each iteration of the estimation, there are many potential combinations of regressors given the hundreds of rows in the observation equation matrices. Future research should examine the feasibility of a variable selection methodology as a means of dimension reduction.

CHAPTER IV

A MONTE CARLO STUDY OF THE EFFECTIVENESS OF THE PC ALGORITHM AT IDENTIFYING THE STRUCTURE OF A VECTOR AUTOREGRESSION

The modeling revolution started by Sims' (1980) vector autoregression (VAR) fundamentally altered time series analysis. By the time Sims' article was published, studies had established univariate, non-structural time series models were outperforming, at least in terms of out-of-sample forecasting, the large macroeconometric models of the 1970s (Nelson 1972; Cooper and Nelson 1975; Cicarelli and Narayan 1980). Univariate models, however, generally cannot provide information necessary for policy implementation. By including multiple time series, the VAR explicitly captures relationships among variables. One major drawback of the VAR, however, is the curse of dimensionality. As additional variables are added to the system, the number of parameters to be estimated in the unrestricted VAR system rapidly increases, decreasing degrees of freedom and increasing the uncertainty surrounding the underlying parameters (Bruggemann 2004).

VAR modeling approaches addressing the curse of dimensionality have been developed. These approaches include Bayesian methods (Litterman 1986), subset VARs (Hsiao 1979), and, more recently, factor VARs (Bernanke, Boivin, and Eliasz 2005). Each method imposes restrictions on the parameter space. That is, if the value, distribution, or some prior belief about a parameter is known a priori, its estimation can be simplified.

Hsiao (1979) proposes a subset VAR methodology which Lutkepohl (2005, p. 211) classifies as a type of "sequential elimination of regressors." Most applications of Hsiao's method find the combination of regressors that minimizes an information criterion. Akleman, Bessler, and Burton (1999) propose the use of directed acyclic graphs (DAGs) to select lags in a subset VAR. The primary benefit of the graph-theoretical version of the Hsiao search is not needing an "…a priori ranking of the importance of lags of each 'other variable'" (Akleman, Bessler, and Burton, 1999 p. 518). In other words, all lags of all variables are determined simultaneously using an algorithm of inductive causation.

Demiralp and Hoover (2003) evaluate the effectiveness of the PC Algorithm for identifying the contemporaneous structure of the VAR, that is, the placement of zeros and ordering of the elements of the contemporaneous relationships between variables. I, however, am unaware of a Monte Carlo study on the PC Algorithm's ability to detect the structure of the lag coefficients in a VAR. The objective of this research is to investigate the appropriateness of the PC Algorithm as a subset VAR methodology in determining both the contemporaneous and lag structure of the data-generating process (DGP). Specifically, the Akleman, Bessler, and Burton (1999) DAG-based methodology using the PC Algorithm (Spirtes, Glymour, and Scheines 2000) is evaluated using Monte Carlo simulation following Demiralp and Hoover (2003). Two results designed to aid applied analyses are provided. First, significance levels for the PC Algorithm's search and discovery of the structure are suggested. Second, the PC Algorithm is assessed for its ability to identify the structural innovations of a VAR.

The Problem of Zero Coefficients

Following textbook treatment such as Hamilton (1994) or Lutkepohl (2005), the VAR is

$$A(L)z_t = e_t, \tag{4.1}$$

where $A(L)$ is a conformable lag polynomial matrix of degree p, z_t is an m-vector of time series, and e_t is a vector of innovations of the same dimension as z_t. The first term of the lag polynomial matrix, A_0, is a lower-triangular matrix with a diagonal of ones and determines the contemporaneous relationships between the series in z_t. Let the covariance matrix of the innovations, $E(e_t e_t') = \Omega$, be diagonal. This is a structural VAR (SVAR) in which the innovations are orthogonalized and independent. If the covariance matrix is not diagonal, a shock in one series would be accompanied by shocks in other correlated series, and interpretation of the "tangled" responses is problematic. Empirical research has been primarily concerned with the placement of zeros in A_0, the matrix designating the contemporaneous relationships between the time series. Once the structure of A_0 is known, the transformation

$$\begin{aligned} A_0^{-1}A(L)z_t &= A_0^{-1}e_t \\ B(L)z_t &= w_t, \end{aligned} \tag{4.2}$$

where $B(L) = A_0^{-1}A(L)$ and $A_0^{-1}e_t = w_t$, provides the reduced-form representation of the VAR in which the innovations are likely correlated, that is, $E(w_t w_t')$ is most likely not diagonal. The reduced-form representation is the most commonly used form in estimating VARs.

The structure of $A(L)$, however, is unknown and has the potential for zero coefficients in its both its contemporaneous and lag matrices. If a variable does not

affect some variable(s) in the VAR, including itself, at a given lag, the coefficient on that value is zero. The ability to impose zero restrictions prior to estimation reduces the dimensionality of the estimation problem. The difficulty, of course, is obtaining a priori knowledge of the location of zero restrictions in $A(L)$. This is the problem for which methodologies addressing the curse of dimensionality have been developed.

Hsiao's (1979, 1981) search methodology for a bivariate vector of stationary time series x_t and y_t begins by fitting univariate autoregressive processes to the individual series by minimizing Akaike's Final Prediction Error (FPE) to determine the optimal lag length. Next, a model is constructed for y_t as a function of its own lags (the number of which is suggested by the univariate modeling procedure). Next, lags of x_t are sequentially added, and as an additional lag is added, the FPE is calculated, and the number of lags of x_t included is determined again by the combination of lags of x_t and y_t that minimizes the FPE. If the FPE of the best model including lags of x_t is greater than that of the univariate models, the univariate process is selected for y_t. The procedure is repeated to determine the order of the lags for x_t. Once the lag structures are determined, the equations are reestimated using a systems estimator to account for the cross-equation correlation of the errors. Hsaio (1979, p. 555) recommends diagnostic checks on the residuals of the system "Because the sequential procedure may bias the joint nature of the process and the single-equation approach is equivalent to ignoring the effect of possible correlations within the components of innovations…" Applications of the Hsiao methodology include Ahking and Miller (1985), Bessler and Babula (1987),

Karfakis and Moschos (1990), Ghatak, Milner, and Utkulu (1997), Liu, Song, and Romilly (1997), and Tan et al. (2012).

Most studies use an information criterion-based approach as initially proposed by Hsiao (1979, 1981). Akleman, Bessler, and Burton (1999, p. 509), however, "…offer [DAGs] as an alternative to regression-based search procedures for placing zero restrictions on relationships among a set of variables." Their procedure uses graph theory and algorithms of inductive causation to determine the appropriate lags. They study corn exports using three-variable VARs using Hsiao-like subset VAR methods, including Hsiao's regression procedure using the FPE, Hsiao's regression procedure using the Schwarz Information Criterion (SIC), and finally a regression-free DAG approach, to determine zero restrictions. Mean squared errors (MSE) of out-of-sample forecasts are generally smaller when using the DAG or SIC to determine the VAR's restrictions. For only one forecast did the MSE of the model suggested by the FPE outperform the models determined using either the SIC or DAG method.

Bruggeman (2004) classifies subset VAR procedures into two categories, single-equation and system. In single-equation techniques, regressors are sequentially deleted equation by equation. These techniques are preferable when there are no contemporaneous relationships between the series. System-based searches typically incorporate a search of all possible variables in every equation. If there is contemporaneous causality amongst the innovations, a systems search is preferable. In a Monte Carlo evaluation of subset VAR methodologies, Bruggeman (2004) finds many of the subset methodologies correctly discover the underlying data-generating process.

He did not, however, evaluate the DAG approach nor Hsiao's method. Simple and complex search procedures are about equally successful. In addition, the forecasting performance of subset VARs was often superior to that of unrestricted VARs.

DAGs and the PC Algorithm

A graph is an ordered triple of vertices *V*, marks *M*, and ordered pairs *E* between the vertices, where *V* and *M* are non-empty sets (Spirtes, Glymour, and Scheines 2000). In the VAR application, vertices represent each time series at each included lag. The ordered pairs *E* describe the position of edges between vertices, and the presence and direction of arrows denote the flow of information. A directed path between two vertices *X* and *Y* occurs when there is a unidirectional sequence of arrows between variables beginning at source *X* and ending at *Y*. Paths which pass through each of their vertices only once are acyclic. A DAG contains no cyclic paths between vertices. DAGs are commonly used to describe the transmission of information among a set of variables.

DAGs are a convenient way of predicting contemporaneous or intertemporal relationships between variables. Algorithms of inductive causation provide a means of estimating these structures. As in any statistical problem, unfortunately, the researcher does not know the true causal relationships, so assumptions regarding the data and the underlying DGP are necessary so the structure can be empirically estimated.

Three major assumptions of the PC Algorithm are the faithfulness condition, the Markov condition, and causal sufficiency (Spirtes, Glymour, and Scheines 2000). Under faithfulness, the lack of an edge between two variables in a graph is the result of no true relationship between the variables and not of cancellation of underlying structural

parameters. The Markov condition states the joint probability of a graph can be written as the product of all events conditioned on only the variable(s) by which they are immediately preceded. The final and most tenuous assumption is causal sufficiency. This assumption requires there are no omitted causes of two or more variables in the set of variables considered. In many VARs, the set of variables considered is quite small because of the curse of dimensionality. The small number of endogenous series suggests that there is a high potential for violations of the causal sufficiency assumption. In addition, it is assumed the series are Normally distributed.

Spirtes, Glymour, and Scheines (2000, p. 84-85) provide a thorough discussion of the algorithm which is summarized here. The algorithm begins with a complete graph of adjacent, undirected edges connecting every variable in the set *V*. Two nodes are said to be adjacent if there is an edge connecting them. To begin, let $n = 0$.

1. Test for (conditional) independence of the variables.
    a. Two adjacent nodes *x* and *y* in *V* (an ordered pair) are selected such that the set of adjacent nodes to *x* in *V* but not including *y* has cardinality (number of elements in the set) of at least *n*. A subset *S* of these adjacencies not including *y* and of cardinality *n* is selected. The variables *x* and *y* are tested for independence conditional on the nodes in *S*. When $n = 0,$ this is equivalent to calculating the conventional unconditional correlation between the series. If the two are independent as deemed by the test statistic and the specified significance level (discussed later), the

edge between *x* and *y* is removed, and the subset *S* is recorded as a member of the sepset between *x* and *y* and between *y* and *x*.

b. Step (1a) is repeated until all possible ordered pairs of nodes in *V* satisfying the criteria in (1a) have been tested for (conditional) independence. The value of *n* is then incremented, and the testing continues by starting over at step (1a). When all sets of adjacencies in *V* for the ordered pair *x*, *y* not including *y* have cardinality less than *n*, step (1) of the algorithm is complete.

2. After edges have been removed, the skeleton is in place. The edges must now be ordered.

   a. Select from *V* nodes *x*, *y*, and *z* such that *x* and *y* are adjacent, *y* and *z* are adjacent, but *x* and *z* are not adjacent. If the node *y* is not contained in the sepset of *x* and *z*, then orient the collider $x \rightarrow y \leftarrow z$. Repeat step (2a) for all ordered triples satisfying these criteria.

   b. For *x*, *y*, and *z* in *V*, if 1) $x \rightarrow y$, 2) *y* and *z* are adjacent, 3) *x* and *z* are not adjacent, and 4) not $z \rightarrow y$, then $y \rightarrow z$. If there is an edge between *x* and *y* and a directed path from *x* to *y*, then $x \rightarrow y$. Repeat step (2b) until finished.

To test statistically the (conditional) independence of two series, a hypothesis test of the series' (conditional) correlation ρ is performed with null and alternative hypotheses $\rho = 0$ and $\rho \neq 0$ (Spirtes, Glymour, and Scheines 2000, p. 94). Fisher's Z transformation of the sample (conditional) correlation $\hat{\rho}$ is

$$z = \frac{1}{2}\ln\left[\frac{1+\hat{\rho}}{1-\hat{\rho}}\right]$$

and is Normally distributed with mean $\frac{1}{2}\ln\left[\frac{1+\rho}{1-\rho}\right] = \frac{1}{2}\ln\left[\frac{1+0}{1-0}\right] = 0$ and standard error $1/\sqrt{N-|S|-3}$, where $|S|$ is the cardinality of the conditioning set *S* in the calculation of the (conditional) correlation from step (1a). The standardized test statistic for the hypothesis test is

$$\frac{\sqrt{N-|S|-3}}{2}\ln\left[\frac{1+\hat{\rho}}{1-\hat{\rho}}\right].$$

If the corresponding p-value is smaller than the specified significance level, there is statistical evidence of a nonzero (conditional) correlated between the series. In the case when $n = 0$ in the algorithm, then $|S| = 0$.

Methodology

To determine the appropriateness of the PC Algorithm at placing zero restrictions in the coefficient matrices of a VAR process, a Monte Carlo experiment is performed. In this experiment, known contemporaneous and lagged structures are used to simulate datasets with various levels of innovation noise. These datasets are then evaluated by the PC Algorithm, and the selected graph is compared to that of the true process to measure the algorithm's success.

Simulation

Various VARs of either two and three variables with either one or two lags are simulated. Let *p* be the maximum number of lags in the true underlying DGP. In each

iteration of the simulation, the coefficients of $A(L)$, including the placement of zeros, change. In other words, the structure of $A(L)$ changes to allow the strengths of the lagged and contemporaneous relationships to vary from none – a zero coefficient – to relatively strong. Demiralp and Hoover (2003) and Bryant, Bessler, and Haigh (2009) employ multiple levels of "relationship strengths" to test the effect of the strength of the relationship on the algorithm's efficacy, but each additional level necessitates many additional simulation iterations.

For each of the p lag matrices in $A(L)$, let a_{lij} denote the coefficient at the ith row and jth column of the lth lag matrix. In this simulation, each a_{lij} can have zero effect or it can have nonzero effect in the open interval -1 to 1. In the case of a nonzero effect, a coefficient is drawn uniformly between zero and one, not inclusive and transformed to be negative with probability 0.5. For example, a two-variable, one-lag SVAR in which x_t is a function of itself and y_t is a function of itself and x_t is

$$\begin{bmatrix} 1 & 0 \\ a_{021} & 1 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} a_{111} & a_{112} \\ 0 & a_{122} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{bmatrix}, \tag{4.3}$$

where $\begin{bmatrix} x_t & y_t \end{bmatrix}' = z_t$ is the vector of time series at time t, $\begin{bmatrix} \varepsilon_{1,t} & \varepsilon_{2,t} \end{bmatrix}' = e_t$ is a vector of structural innovations at the same t. The "0" represents a zero coefficient in the true DGP, and the other a_{lij} are assumed to be nonzero. For this experiment, every combination of zero and non-zero placements is explored. The graphical representation of this DGP is shown in Figure 4.1. Corresponding to the coefficient placement in equation (4.3), the arrows in the figure show information flows from lags of both variables to x_t and from y_{t-1} to y_t. The flow of information from x_t to y_t corresponds to

a_{021}, the coefficient in the contemporaneous design matrix. In the discussion below, the a_{lij} refer to the lag structure in the experiment when $l >= 1$ and to the contemporaneous structure of the experiment when $l = 0$.

After the SVAR coefficients have been drawn, the VAR is transformed into its reduced form as in equation (4.2) to simplify the simulation. The reduced-form VAR is tested for stationarity by determining the roots of the polynomial

$$\left| I - B_1 r - B_2 r^2 - \ldots - B_p r^p \right|, \tag{4.4}$$

where I is the identity matrix, B_k is the kth member of the lag polynomial matrix $B(L)$, r is a conformable vector, and $|.|$ is the matrix determinant. If the determinant is nonzero and its roots are outside the unit circle, then the VAR is stationary (Lutkepohl 2005). If the VAR is nonstationary, the SVAR coefficients are redrawn until a stationary DGP is obtained.

Once the DGP is determined, observations are simulated. For each structural VAR, two values, one and four, are used for the standard deviation of the white noise, which is assumed to be Normally distributed. This vector of uncorrelated white noise is premultiplied by A_0^{-1} as in equation (4.2) to obtain the correlated innovations. The starting vector for each simulated series is the zero vector; a warm-up period of 1,000 observations is used to mitigate starting value effects. Each observation is generated by the sum of its deterministic components (the products of the lags and their coefficients) and the stochastic component (correlated residuals). An additional, informal test of stationarity is performed on the simulated series. The series is simulated again if the

difference between the starting and ending values of the chain is greater in magnitude than the product of eight and the residual standard deviation.

To use the PC Algorithm to determine the lag structure, the researcher must include one variable for each lag believed to affect the contemporaneous values of each series. The researcher, of course, has no knowledge of the true number of lags, so it is necessary to perform a search over possible lags. In the simulation, for each dataset, the graph is estimated under several search lags to imitate a researcher's lack of knowledge regarding the lag structure. The PC Algorithm is evaluated with between one and four lags for each test.

Evaluation

The estimated DAG is evaluated relative to the true DGP. Using the language of Demiralp and Hoover (2003), the PC Algorithm-estimated graph is called the selected graph, and the true graph defined by the actual structure of the DGP is the reference graph. A correct link occurs when the selected graph exactly matches the reference graph on a given arrow. An edge is unresolved when it is directed in the reference graph but no orientated in the selected graph. An error of omission occurs when the entire edge is not included in the selected graph but is present in the reference graph. This occurs when the PC Algorithm suggests no relationship between two nodes when there is a relationship between the variables. An insertion occurs when the selected graph includes an edge that is not in the reference graph. Finally, a reversal error occurs when a mark on the selected graph occurs on the opposite end of an edge in the reference graph. Because the DGP is a time series, the only possible way of obtaining a reversal

error is when evaluating the selected graph's suggested contemporaneous structure. Arrows, in theory, cannot be oriented from future to past, but in practice, however, reversals do occur. One way to prevent the PC Algorithm from orienting from future to past is by providing background knowledge preventing backwards orientations. This requires slight modification to the PC Algorithm as outlined in Spirtes, Glymour, and Scheines (2000, p. 93). This modification is applied here.

Following Tsamardinos et al. (2006) and Kalisch and Buhlmann (2007), the performance metric utilized is the Structural Hamming Distance (SHD), which is the count of omitted, reversed, and inserted edges in the selected graph relative to the reference graph. Other things equal, a smaller SHD is preferred and implies fewer errors are made in determining the true structure of the DGP and suggests the PC Algorithm could be used in a subset VAR procedure.

A graphical depiction of each potential error for the DGP given in equation (4.3) is shown in Figure 4.2. The top graph is the reference graph, the true graph described by the DGP; this graph matches that of Figure 4.1. In the leftmost pane of the second row, an unresolved edge is shown between x_t and y_t because, while the algorithm correctly identified the edge, it was unable to determine in which direction the causality flows. In the second pane of the second row, an error of omission occurs because the directed edge from y_{t-1} to x_t is missing. An insertion is demonstrated in the third pane as an edge is added from x_{t-1} to y_t. In the final pane, a reversal is shown because the algorithm identified the information flowing backwards from y_t to x_t.

SHD metrics are calculated separately for the lag structure, the contemporaneous structure, and finally the total of both the lag and contemporaneous SHD valaues. For the second, third, and fourth graphs shown in the second row of Figure 4.2, the value of the total SHD considering both lag and contemporaneous structures is one. The total SHD of the first graph in the second row is zero because unresolved edges are not penalized by this metric. The values of the SHD for the lag portion of the DGP are zero (no errors considered by SHD) in the first and last graphs and one in the second and third graphs (an omission or insertion). The contemporaneous SHD values are zero for all except the last graph where it is one.

When interpreting SHD values across the different DGPs, it is important to consider the number of relationships in the DGP. One relationship corresponds to one coefficient in the VAR structure, either lagged or contemporaneous. The number of nodes in the reference graph is the product of the number of series and one plus the number of lags. There are, however, potentially more relationships than nodes because there can be an edge from any node to any other node that does not violate the time ordering of the nodes. To clarify, consider the DGP represented by the DAG in Figure 4.1. There are four total nodes considering periods t and $t-1$ (one for each x and y). Arrows could potentially occur, however, from each node at $t-1$ to each node at t (four lag relationships in total). There could potentially be an arrow between x_t and y_t (one contemporaneous relationship). A summary of the number of nodes and relationships for the DGPs in this experiment follows.

- Two-variable, one-lag VAR (2V1L): There are four nodes in the reference graph which gives four potential lag relationships and one potential contemporaneous relationship.
- Two-variable, two-lag VAR (2V2L): There are six nodes in the reference graph which gives eight potential lag relationships and one potential contemporaneous relationship.
- The three-variable, one-lag VAR (3V1L): There are six nodes in the reference graph which gives nine potential lag relationships and three contemporaneous potential contemporaneous relationships.

To simplify comparisons across DGPs, values of the SHD are normalized by the number of potential relationships in the reference graph, referred to as the SHD per relationship. As in the case of the SHD discussed previously, three SHD per relationship metrics are presented. First, the SHD per lag relationship (SHDLR) considers only the lag portion of the graph only. Next, the SHD per contemporaneous relationship (SHDCR) considers only errors in the contemporaneous design matrix. Finally, the SHD per the total number of relationships (SHDTR) measures the total error per relationship in the graph. The SHDTR is not the sum of the SHDLR and SHDCR because their denominators differ.

Summary

The entire procedure is repeated to accommodate combinations of sample size, significance levels, and the number of lags included in the PC Algorithm's search. To summarize, the procedure is

1. Specify the lag and contemporaneous structure of the VAR to be tested. Each lag and contemporaneous relationship coefficient is drawn according to the current structure which is either zero or in the range between -1 and 1 not inclusive and not including zero. Test each VAR structure for stationarity. If the process is not stationary, redraw coefficients for this structure.
2. For each level of innovation white noise, simulate a vector time series according to the DGP created in the first step. A warm-up period of 1,000 observations is discarded to mitigate the effects of the starting values (a vector of zeros). Perform a second, informal test for stationarity of the time series by calculating the difference between the starting and ending values. Any series with a difference greater than a threshold (the innovation standard deviation multiplied by eight) prompts a redraw of the time series.
3. After obtaining a stationary set of time series, the PC Algorithm is used to test the VAR under each combination of the following scenarios:
    a. Number of observations: 50, 100, 250, 500, and 1,000 (five total);
    b. Significance levels: from one to 20 percent in increments of one percentage point, plus 0.5, 25, and 30 percent (23 total); and
    c. Number of lags tested in the PC Algorithm: one through four (four total).

4. Evaluate the selected graph against the reference graph for errors of omission, reversal, addition, as well as edges that are unresolved.

For each simulated dataset, then, the PC Algorithm is evaluated for a total of $5 \times 23 \times 4 = 460$ times or

number of observations $\times$
number of significance levels $\times$
number of search lags.

In addition, for each structure, 100 datasets are drawn for the DGPs with two variables; 25 datasets are drawn for the DGPs with three variables. The reduction in the number of datasets generated for the two- and three-variable DGPs because of the increased computational time required to simulate and analyze the results of the three-variable VAR. The total number of iterations is

number of contemporaneous structures $\times$
number of lag structures $\times$
number of residual variance levels $\times$
number of evaluations $\times$
number of datasets,

which for the 3V1L, $2^3 \times 2^9 \times 2 \times 460 \times 25 = 94,208,000$. Similarly, the number of simulations for 2V1L is 2,944,000, and the number of simulations for 2V2L is 47,104,000.

The experiment is written in C utilizing free and open source software including OpenMPI (The OpenMPI Project 2018), OpenMP (OpenMP 2018), the GNU Scientific Library (Free Software Foundation, Inc. 2018), OpenBLAS (Xianyi 2018), and dSFMT (Saito and Matsumoto 2013). Java code for the PC Algorithm and Tetrad are available

from the Tetrad Project (2017) and adapted for use here. The experiments are performed using high performance computing resources at Texas A&M University.

System R-Squared

For each simulated system, the system R-squared (Vahid and Issler 2002), an extension of the traditional R^2 to a VAR system is calculated to determine the relative strength of the system's coefficients compared to the other simulated systems. The system R-squared is

$$\frac{trace(\Gamma_0\Omega^{-1}) - N}{1 + trace(\Gamma_0\Omega^{-1}) - N}, \tag{4.5}$$

where $\Gamma_0 = E(z_t z_t')$, $\Omega = E(\varepsilon_t \varepsilon_t')$, $E(\cdot)$ is the expectation operator, and the other variables are as previously defined. The population expectations needed to calculate equation (4.5) can be calculated in a straightforward manner because the DGP is known and stationary.

Results

The proportion of graphs with an SHDTR of zero (no errors) are shown by DGP in Table 4.1. The additional complexity of the 2V2L and 3V1L DGPs compared to the 2V1L result in a substantial reduction in the proportion of graphs with zero SHDTR.

Evaluating the Lag Structure of the DGP

A graphical summary of the simulation results for the lag portion of the DGP is presented in Figure 4.3. The average SHDLR across all datasets for every combination of DGP, the number of variables and lags in the DGP and the standard deviation of the residuals, is plotted as a function of the number of lags searched in the PC Algorithm

search, the number of observations, and the significance level used in the PC Algorithm's hypothesis tests. Several patterns emerge from this broad view of the simulations: 1) robustness to residual variability, 2) the benefits of additional observations, and 3) the choice of significance level and the number of search lags included substantially affect search accuracy.

Variability of the Residuals. First, comparing across the same variable and lag combination but with different levels of residual volatility (for example, the top-left panel in Figure 4.3 compared to the panel below it), the results are almost identical. The SHDLR generally is not affected by a fourfold increase in the noise of the process for the DGPs tested. This suggests the detection of the VAR's lag structure is fairly robust to the innovation white noise. To simplify the presentation, the results presented in Figure 4.3 are re-presented in Figure 4.4, where the partitioning of the effects by standard deviation of the residuals is removed.

Number of Observations. In Figure 4.4, the advantage of having additional observations is clear. Increasing the number of observations consistently decreases the SHDLR for a given significance level. In Table 4.2, the average SHDLR is shown by the number of observations in the series. For all three structures tested, the benefits of 1,000 observations versus 50 observations in determining the lag structure are approximately a 30 to 40 percent reduction in the SHDLR. The benefits, however, are not directly proportional to the number of additional observations. When moving from 50 to 100 observations, a doubling of observations, the reduction in the average SHDLR is between 11 and 16 percent, but an additional 500 observations in a dataset with 500

observations already, another doubling of observations, only reduces the average SHD up to eight percent.

Significance Level Selection in PC Algorithm. The SHDLR tends to be minimized when the significance level used in the PC Algorithm is small. In Figure 4.5, the results are presented where the partitioning of the effects by lags used in the search is removed. In applied analysis, researchers likely only have accurate knowledge of the number of observations available for estimation. In most cases, researchers know which variables might belong in the system, but there is still uncertainty, and the existence of omitted variables is likely. As such, Figure 4.5 is useful for determining the "optimal" significance level to utilize in the PC Algorithm when constructing a subset VAR because its level of detail reflects the researcher's lack of knowledge of the true number of lags or variables in the DGP. Selection of an appropriate significance level substantially improves the PC Algorithm's discovery of the lag structure. Depending on the number of observations, the SHDLR is reduced between 18 and 60 percent by moving from the significance level that performed the worst to the level that performed best. For all DGPs examined, the two percent significance level generally provides the smallest SHDLR. In the case of datasets with 50 observations, the optimal significance level is five percent. Using the smallest significance level tested, 0.5 percent, yields a notable increase in the SHDLR over the two percent level.

A major exception to the pattern regarding the choice of significance level discussed previously is identified in Figure 4.4. When the number of search lags used in the PC Algorithm exactly equals the number of lags in the DGP, the two percent

significance level does not appear to be optimal for smaller datasets. For datasets with 50 observations, significance levels from 15 to 30 percent are all almost equally effective at identifying the lag structure of the VAR. For datasets with 100 observations, significance levels from 10 to 20 percent are almost equally effective. Given a small dataset with prior knowledge of the number of lags in the DGP, a practitioner should be equally as effective at discovering the lag structure of the VAR using any significance level in these ranges. In most cases, however, such knowledge is not available, and the findings generalized from Figure 4.5 should be used to guide applied analyses.

Lag Selection in PC Algorithm. In VAR analysis, careful selection of the number of lags is important to capture the dynamics of the system. As discussed previously, when constructing a graphical version of the VAR using an algorithm of inductive causality, the researcher is also faced with the selection of the lags to search. In practice, this choice involves a balance between under- and over-fitting, or the inclusion of insufficient or too many lags compared to the true number of lags in the DGP. As expected for all DGPs, the SHDLR is minimized when the number of search lags exactly equals the number of number of lags in the DGP (Table 4.3). That is, if searching over the true model (the correct variables and the correct number of lags), the number of errors is minimized. This minimization is expected because one cannot obtain a graph with too few or too many lags. While it does increase the SHDLR, the inclusion of extra lags, however, increases the SHDLR less than the inclusion of too few lags. In the first panel of the second row in Figure 4.4 with 2V2L structures tested using only one lag in the PC Algorithm, the SHDLR is larger than in the following three frames. In fact, the

benefit of small significance levels in the case of insufficient lags is only marginal. The SHDLR for graphs estimated with insufficient lags is on average 48 percent larger than the SHDLR for graphs with at least the number of lags in the DGP for the simulated 2V2L. The effect of including extra lags is shown in Figure 4.6. The penalty incurred by the inclusion of extra lags is partially offset by using small significance levels. In the case of the 2V1L with 250 observations evaluated at the two percent level, the SHDLR increases 41 percent when extra lags are included. Using a ten percent significance level, the SHDLR increases 96 percent. At the largest significance levels tested, the SHDLR increases considerably for graphs constructed using extra lagged values. When unsure of the correct number of lagged values to include, a smaller significance level should be utilized to minimize the number of errors.

Overview of Findings for the Lag Portion of the DGP. From the preceding, the following inferences arise when implementing the PC Algorithm to obtain the SHDLR-minimizing graph for a VAR's lag structure with no prior knowledge of the DGP.

1. A two percent significance level is optimal in many, but not all, cases.
2. Over-fitting in the number of search lags in the PC Algorithm produces smaller SHDLRs compared to not including sufficient lags. While there is a penalty for including lags exceeding the true number of lags, the penalty is decreased by using smaller significant levels.
3. Additional observations are beneficial but maybe not always possible based on data limitations.

Evaluating the Contemporaneous Structure of the DGP

The average SHDCR for every combination of data structure plotted in Figure 4.7 as a function of the number of lags included in the PC Algorithm search, the number of observations, and the significance level used in the PC Algorithm's hypothesis tests. Similar patterns to those seen in Figure 4.2 regarding the number of observations and the significance level are present in Figure 4.7. Datasets with more observations have smaller SHDCR values. Employing a small significance level decreases the SHDCR.

The average SHDCR is shown by number of observations in Table 4.4. For each DGP, the SHDCR decreases as more observations are available. For the 2V1L inspected, there are insertions, deletions, and reversals in 20 to 33 percent of the graphs evaluated, decreasing as the number of observations in the dataset increase. These values for the contemporaneous portion of the graph are similar to their lag-portion counterparts. Similar findings are found for the other two DGPs. In all cases, however, the SHDCR is larger – up to 35 percent – than the corresponding values for the lag portion.

Although the SHDCR only evaluates the PC Algorithm's efficacy at determining the contemporaneous relationships between the VAR's variables, knowledge of the true number of lags in the DGP still generally produces a smaller SHDCR. Compared to those of the first, third, and final panes of the same row, the SHDCR values in second pane in the second row of Figure 4.7 are smaller for datasets of at least 250 observations. There is a small, yet consistent penalty to using the incorrect number of lags in the search. In addition, the values of the SHDCR for the 2V1L are slightly smaller, other

things equal, than those of the 2V2L despite both DGPs containing the one contemporaneous relationship. This suggests additional information flowing from the additional lagged variables increase the chances of errors in discovering contemporaneous relationships.

System R-Squared

In Figure 4.8, the SHDTR is plotted against the system R-squared by DGP. A histogram of all simulated datasets with system R-squared values falling into bins each five percentage points wide is also shown. The distribution of system R-squared values provided by the simulated VARs are similar for the 2V2L and 3V1L, but distribution of the 2V1L is different. The discrepancy suggests comparisons across DGPs may not be valid as the larger DGPs generally have higher values of the system R-squared.

In most cases, the SHDTR is minimized when the system R-squared falls between 20 and 60 percent. For systems with small R-squared values, there is likely insufficient information to produce the correct graph, while for systems with larger R-squared values, there is likely too much information to produce the correct graph. While the population system R-squared is unknown to the researcher, it provides insight here into the kinds of VAR systems the PC Algorithm is best able to identify.

Discussion and Conclusions

The PC Algorithm is examined for its ability to propose zero restrictions on the lag and contemporaneous matrices in a VAR for different sizes of datasets, different combinations lagged values, and significance levels. This extends the work of Demiralp and Hoover (2003), who evaluate the PC Algorithm's ability to detect the causal

ordering of a VAR postestimation using its residuals. The experiments suggest the PC Algorithm can be effective at discovering the lag structure of a VAR. In addition, it appears to be almost equally effective at obtaining the contemporaneous portion of the DGP and the lag portion. In both cases, the efficacy is increased when more observations are available to the researcher and the appropriate number of lags is employed in the search graph. Obtaining more data is often impractical, and researchers do not have prior knowledge of the appropriate number of lags to include in the search. If they did, there would be no reason to search for the number of lags. As such, results suggest significance levels between two and five percent should be used in the hypothesis testing portion of the algorithm regardless of the number of observations.

The subset VAR methodology is a potential solution for VAR analysis of larger datasets. This methodology, however, can be limited in its efficacy or be rendered ineffective if the DGP implies no zero restrictions. In these cases, the subset VAR methodology is unable to mitigate the effects of the curse of dimensionality, and another solution for dimension reduction is necessary.

For DGPs for which the subset VAR methodology can assist in the dimensionality problem, there are still other prominent issues in practice. Sufficient data and knowledge of the lag structure are likely not the only problems faced when employing the PC Algorithm for causal discovery. Given the causal sufficiency assumption of the PC Algorithm, mistakenly assuming the incorrect variables to include in the DGP will likely lead to additional errors regardless of the search lag and significance level used. For the presented experiments, the DGPs were all stationary,

Gaussian processes, but non-Gaussian data and nonstationary time series are both common in applied research. While the PC Algorithm is unable to infer causality for these types of data, other algorithms are available. Further research should provide Monte Carlo evidence of their efficacy in the construction of subset VAR models when the assumptions of the PC Algorithm do not hold. As also suggested by Demiralp and Hoover (2003), the problem of selecting the maximum number of lags to be included in the PC Algorithm is still not solved. The results here suggest over-fitting, while prone to added insertion errors, produces a smaller SHDR on average than under-fitting.

CHAPTER V

CONCLUSION

Manifestations of the curse of dimensionality have changed. What once was a lack of computing resources to handle high-dimensional problems is, at least in applied time series econometrics, a potential need for new methodologies to handle the larger and larger datasets commonly available for analysis. In this dissertation, three studies are presented that were affected by the curse of dimensionality.

In the ductless heat pump (DHP) adoption study, the curse is addressed, as most studies do, by ignoring many potential variables. Adoption is modeled using an extension (Jain and Rao 1990) of Bass's (1969) simple model of adoption. This modeling technique is in line with a statement made by Oscar Burt during the 1980s; the essence of modeling is to capture reality as simply as possible. The second study uses Bayesian techniques to include large datasets in the factor-augmented vector autoregression (FAVAR) framework with time-varying parameters. In this methodology, the data included is not limited a priori but rather is used directly in the dimensionality reduction to obtain one or more unobservable factors that evolve over time and explain some common co-movements of the data. Here, the curse is still realized in the computational burdens of estimating and comparing several specifications of a complex model. Through Monte Carlo experiments, the final study evaluates directed acyclic graphs (DAGs) as a potential tool for dimension reduction; data-driven techniques are used to identify the "best" model in line with Burt's statement. Together,

these studies address the overall objective of demonstrating three very different methodologies and applications that all attempt to do the same thing: simplify the not-so-simple.

Summary of Results

Adoption of DHPs

The aim is increasing the understanding of DHP adoption in the Pacific Northwest of the US by quantifying the effect of utility-provided rebates and Northwest Energy Efficiency Alliance (NEEA) expenditures on the number of installations and providing forecasts of DHP installations through 2018 given various rebate and NEEA expenditure levels is addressed by modeling using household-level data provided by NEEA. The adoption mechanism is complex, driven by innovation, communication channels, time, and the social system (Rogers 1962). Each of these is difficult to quantify and, in fact, may be difficult to identify until after adoption has happened. When modeling a complicated socio-economic process, the curse of dimensionality reveals itself as a need for a methodology to capture the adoption profile of consumers that is both flexible for forecasting and identifiable without a large number of variables that, in fact, may not be identified yet. In this case, the simple model of Bass (1969) with extensions by Jain and Rao (1990) are applied. The Bass (1969) model assume a sigmoid-shaped adoption profile that is driven by coefficients of innovation and imitation as well as market potential. The results are directed to policymakers and utility providers in the Pacific Northwest concerned with finding cost-effective methods to promote energy efficiency and reduce future load growth.

Rebates are more effective at increasing installations than are NEEA's expenditures on marketing and installer training because DHP adoptions are elastic with respect to net cost of installation. In addition, the reduction of federal tax incentives after 2011 decreases the probability of adoption. The adoption profile for DHPs grows slowly; consultation with NEEA indicated the forecasted adoption rate is slower than their goal. Three reasons are posited for the slow growth. First, the up-front costs of adoption are high. In addition, DHPs are not "leisure-enhancing," making adoption rates lower. Finally, economic conditions during and after the Great Recession, the period studied, may have inhibited adoptions.

Dynamics of the Oil Prices and the Petroleum Industry

The objective of the second study is to investigate the dynamic effects of shocks in oil supply, aggregate demand, and oil demand on oil prices, the upstream, midstream, and downstream sectors of the petroleum industry, and the broader US economy. The dynamics effects are studied over time in a data-rich environment. Baumeister and Peersman (2013) have similarly studied the effects of oil supply shock using Primiceri's (2005) time-varying parameter (TVP)-VAR, but a similar TVP investigation of these shocks has not been performed in a data-rich environment. Because no single variable completely describes an economy or industry, a large set of informational time series are selected, representing data from the upstream, midstream, and downstream of the petroleum industry, as well as interest rates, money supply, exchange rates, employment and earnings, unemployment, and indices of real activity, production, and prices. The findings are beneficial for decision makers inside the petroleum industry to understand

the magnitude and duration of shocks in oil supply and demand as well as aggregate demand.

The curse of dimensionality is encountered here as analysis of a large set of series using traditional, small-scale VARs is infeasible. As such, the TVP-FAVAR model (Mumtaz, Zabczyk, and Ellis 2011) is employed as a solution to the dimensionality dilemma. This technique addresses the curse by using all the data simultaneously to reducing the co-movements of the data to one or more latent factors. Although feasible, Bayesian estimation of the TVP-FAVAR using Markov Chain Monte Carlo techniques requires considerable computational time and memory. Depending on the model specification, one chain of the Markov Chain Monte Carlo estimation required up to 40 hours. In addition, computation of the Deviance Information Criterion, using the particle filter for estimation of the likelihood, for model comparison requires substantial computing time made possible by parallelization. In this application, several specifications were run simultaneously using advanced computing resources that are not always available to researchers. As such, the methodology is not feasible for all applications.

The effects of oil supply shocks on oil prices, production, and storage increase during recessions. For oil prices, the findings here support those of previous studies showing oil demand shocks generate the largest responses in magnitude of the shocks studied. On the other hand, for crude oil production, shocks in aggregate demand commonly have a larger impact than oil demand shocks. As oil price volatility increases, the responses of oil prices to oil demand also become larger in magnitude.

The dynamics of oil prices and oil storage show the most variability over the last 25 years, while those of field production, refinery capacity, and gasoline sales appear to be relatively time-invariant. Real activity, as measured by the Industrial Production Index, also appears to be relatively time-invariant in its dynamics, but the dynamics of the Producer Price Index are influenced by oil price volatility.

Subset VARs and Directed Acyclic Graphs

Another solution to the curse of dimensionality in VAR analysis is the subset VAR methodology which can potentially reduce the number of parameters in a VAR. A regression-based subset VAR approach for determining parameters to constrain is suggested by Hsaio (1979, 1981), whereas Akleman, Bessler, and Burton (1999) propose directed acyclic graphs (DAGs) as a means of determining the placement of zero restrictions. Of course, one could argue the subset VAR approach is not actually a reduction of dimensionality but rather identification of the model that correctly describes the dynamics of multivariate time series. Identification of the correct model, however, often results in a lower dimension model.

Here, the objective is to investigate the appropriateness of the PC Algorithm as a subset VAR methodology in determining both the contemporaneous and lag structure of the data-generating process (DGP). Previous research (Demiralp and Hoover 2003) has examined the PC Algorithm's ability to identify the placement of zero restrictions in the contemporaneous portion of the VAR structure. To my knowledge, no study has investigated: 1) the placement of zeros for a subset VAR methodology on the lag portion of the DGP; 2) choices in significance level on the graphs recovered for VARs; and 3)

how searching over different numbers of lags impacts the efficacy of the PC Algorithm at identifying the VAR structure. The results are directly applicable to researchers who model using VARs 1) to reduce the parameterization; 2) to identify the "correct" model for the specific DGP; or 3) to improve forecast performance using subset VARs (Akleman, Bessler, and Burton 1999; Bruggeman 2004).

The PC Algorithm is equally effective uncovering the contemporaneous and lag portions of the DGP and, as such, is potentially useful tool for constructing subset VARs. For the DGPs considered, significance levels between two and five percent are optimal for minimizing the value of the Structural Hamming Distance per relationship. Other things equal, the algorithm is also more effective when there are more observations. When researchers do not have a priori knowledge of the true number of lags in the data-generating process, overfitting provides a smaller penalty than underfitting.

Limitations and Suggestions for Further Research

The coefficients of innovation and imitation play a key role in the shape of the adoption function (Bass 1969), and estimation of these parameters using limited data may yield inaccurate forecasts (Heeler and Hustad, 1980). Particularly, the adoption curve is fit best when the peak of instantaneous adoptions has been realized. Now that the forecast window has almost been realized, future research should evaluate the accuracy of the probability forecasts generated from the adoption model. In addition, data were only provided for households receiving financial incentives for adoption, but adoptions of DHPs that did not receive incentives are likely affected by adoptions that did receive

incentives (and vice versa). Analyses quantifying the effects of "cross-adoption" are needed. Finally, the adoption profile for DHPs should be compared to those of other energy-efficient technologies and under different modeling assumptions to examine if and how the adoption profile varies by type of technology.

The complexity of the Bayesian TVP-FAVAR lends itself to limitations which are the result of necessary simplifying assumptions. One limitation is prior selection. In most cases, conjugate priors are chosen for computational simplicity, so further research should investigate less informative priors. A specification with controls for seasonality might be beneficial given the potential importance including up to the 12^{th} autoregressive lag for the latent factor in the VAR. Identification by sign restrictions in the TVP-FAVAR context, too, is not always straightforward. More informational series allow for more flexibility in the placement of restrictions because many different types of series can be used to identify the shocks, but additional restrictions can lead to a decrease in the number of accepted rotation matrices, making innovation accounting difficult. It is also not clear which variables in a closely related group should and should not be restricted. Amir-Ahmadi and Uhlig (2009) consider this issue for a monetary policy shock in the original Bernanke, Boivin, and Eliasz (2005) FAVAR formulation, but further research should investigate the sensitivity of innovation accounting to the placement of restrictions for the petroleum industry TVP-FAVAR.

Assumptions of the PC Algorithm are often questionable in practice. Many time series are not stationary, and many are not Normally distributed. In addition, the causal sufficiency assumption is likely not to hold for VARs. These experiments should be

repeated using other algorithms of inductive causality that relax these assumptions. In line with Demiralp and Hoover (2003), this study reaffirms the selection of the appropriate number of lags to include for the search remains an issue.

Final Words

Has the curse of dimensionality disappeared from applied time series econometrics? No. Will it ever disappear? Probably not. Just as the science of economic and statistical analyses grows and adapts to the present challenges at hand, so too, will the curse of dimensionality likely continue to evolve and be encountered in different ways for very different research problems.

REFERENCES

Aalbers, R., E. van der Heijden, J. Potters, D. van Soest, and H. Vollebergh. 2009. Technology Adoption Subsidies: An Experiment with Managers. *Energy Economics* 31: 431-442.

Aastveit, K.A. 2014. Oil Price Shocks in a Data-Rich Environment. *Energy Economics* 45: 268-279.

Adams, F.G. and J.M. Griffin. 1972. An Economic-Linear Programming Model of the US Petroleum Refining Industry. *Journal of the American Statistical Association* 67: 542-551.

Agerton, M., P. Hartley, K.B. Medlock, and T. Temzelides. 2014. Employment Impacts of Upstream Oil and Gas Investment in the United States. Working paper, James A. Baker III Institute for Public Policy, Houston. Accessed October 11, 2016. Retrieved from http://bakerinstitute.org/media/files/files/5871a63e/CES-pub-EmploymentImpactsofUpstream-082214.pdf

Ahking, F.W. and S.M. Miller. 1985. The Relationship between Government Deficits, Money Growth, and Inflation. *Journal of Macroeconomics* 7: 447-467.

Akleman, D.G., D.A. Bessler, and D. Burton. 1999. Modeling Corn Exports and Exchange Rates with Directed Graphs and Statistical Loss Functions. In C. Glymour and G. Cooper, eds. *Computation, Causation, and Discovery*. Cambridge: AAAI/MIT Press, 497-520.

Akram, Q. 2009. Commodity Prices, Interest Rates and the Dollar. *Energy Economics* 31: 838-851.

Alessi, L., M. Barigozzi, and M. Capasso. 2010. Improved Penalization for Determining the Number of Factors in Approximate Factor Models. *Statistics & Probability Letters* 80: 1806-1813.

American Petroleum Institute. 2013. *Economic Impacts of the Oil and Natural Gas Industry in the US Economy in 2011*. Washington, DC. Accessed October 11, 2016. Retrieved from http://www.api.org/~/media/Files/Policy/Jobs/Economic_impacts_Ong_2011.pdf

Amir-Ahmadi, P. and H. Uhlig. 2009. Measuring the Dynamic Effects of Monetary Policy Shocks: A Bayesian FAVAR Approach with Sign Restriction. Working paper, Department of Economics, Humboldt University of Berlin, Berlin. Accessed October 3, 2018. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=0C6F6EF3380E1C9C134AEDC9EC1EDD77?doi=10.1.1.676.6664&rep=rep1&type=pdf

An, L., Jin, X.Z., and X.M. Ren. 2014. Are the Macroeconomic Effects of Oil Price Shock Symmetric? A Factor-Augmented Vector Autoregressive Approach. *Energy Economics* 45: 217-228.

Anzuini, A., M.J. Lombardi, and P. Pagano. 2010. The Impact of Monetary Policy Shocks on Commodity Prices. Working Paper, European Central Bank, Frankfurt. Accessed October 11, 2016. Retrieved from https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1232.pdf

Avalos, R., T. Fitzgerald, and R.R. Rucker. 2016. Measuring the Effects of Natural Gas Pipeline Constraints on Regional Pricing and Market Integration. *Energy Economics* 60: 217-231.

Bahadir, B. and W.D. Lastrapes. 2015. Emerging Market Economics and the World Interest Rate. *Journal of International Money and Finance* 58: 1-28.

Bai, J. and S. Ng. 2002. Determining the Number of Factors in Approximate Factor Models. *Econometrica* 70: 191-221.

Barsky, R.B. and L. Kilian. 2004. Oil and the Macroeconomy since the 1970s. *Journal of Economic Perspectives* 18: 115-134.

Bass, F.M. 1969. A New Product Growth Model for Consumer Durables. *Management Science* 15: 215-227.

Bass, F.M., T.V. Krishnan, and D.C. Jain. 1994. Why the Bass Model Fits without Decision Variables. *Marketing Science* 13: 203-223.

Baumeister, C. and L. Kilian. 2016. Understanding the Decline in the Price of Oil Since June 2014. *Journal of the Association of Environmental and Resource Economists* 3: 131-158.

Baumeister, C. and L. Kilian. 2017. Lower Oil Prices and the US Economy: Is This Time Different? Research paper, Center for Economic Policy Research, London. Accessed July 29, 2018. Retrived from https://cepr.org/active/publications/discussion_papers/dp.php?dpno=11792

Baumeister, C. and G. Peersman. 2013. Time-Varying Effects of Oil Supply Shocks on the US Economy. *American Journal of Macroeconomics* 5: 1-28.

Berg, A., R. Meyer, and J. Yu. 2004. Deviance Information Criterion for Comparing Stochastic Volatility Models. *Journal of Business & Economic Studies* 22: 107-120.

Bessler, D.A. and R.A. Babula. 1987. Forecasting Wheat Exports: Do Exchange Rates Matter? *Journal of Business & Economic Statistics* 5: 397-406.

Bellman, R.E. 1961. *Adaptive Control Processes: A Guided Tour.* Princeton: Princeton University Press.

Bernanke, B., J. Boivin, and P. Eliasz. 2005. Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. *The Quarterly Journal of Economics* 120: 387-422.

Bernanke, B., M. Gertler, and M. Watson. 1997. Systematic Monetary Policy and the Effects of Oil Price Shocks. *Brookings Papers on Economic Activity* 28: 91-157.

Binder, K.E. 2016. The Past, Present, and Future of the US Electric Power Sector: Examining Regulatory Changes Using Multivariate Time Series Approaches. Doctoral Dissertation, Texas A&M University, College Station.

Binder, K.E., M. Pourahmadi, and J.W. Mjelde. 2018. The Role of Temporal Dependence in Factor Selection: Implications for Forecasting. *Empirical Economics* 1-39.

Breusch, T.S., and A.R. Pagan. 1980. The Lagrange Multiplier Test and its Applications to Model Specification in Econometrics. *Review of Economic Studies* 47: 239-253.

Bruggemann, R. 2004. *Model Reduction Methods for Vector Autoregressive Processes*. Berlin: Springer.

Bryant, H.L., D.A. Bessler, and M.S. Haigh. 2009. Disproving Causal Relationships Using Observational Data. *Oxford Bulletin of Economics and Statistics* 71: 357-374.

Bugbee, J.E. and J.R. Swift. 2013. Cold Climate Ductless Heat Pump Performance. *Energy Engineering* 110: 47-57.

Carter, C.K. and R. Kohn. 1994. On Gibbs Sampling for State Space Models. *Biometrika* 81: 541-553.

Chen, F. and S.C. Linn. 2017. Investment and Operating Choice: Oil and Natural Gas Futures Prices and Drilling Activity. *Energy Economics* 66: 54-68.

Cheng, B.S. 1998. Oil Prices and Drilling Activity in the United States: An Application of Cointegration and Error-Correction Modeling. *Energy Sources* 20: 459-464.

Chesnes, M. 2009. Capacity and Utilization Choice in the US Oil Refining Industry. Paper presented at 7th International Industrial Organization Conference, Northeastern University, Boston, April 3-5. Accessed October 11, 2016. Retrieved from https://editorialexpress.com/conference/IIOC2009/program/IIOC2009.html#50

Chesnes, M. 2015. The Impact of Outages on Prices and Investment in the US Oil Refining Industry. *Energy Economics* 50: 324-336.

Cicarelli, J. and J. Narayan. 1980. The Performance of Eleven Economic Forecasting Models in the 1970s. *Business Economics* 15: 12-16.

Cogley, T. and T.J. Sargent. 2005. Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US. *Review of Economic Dynamics* 8: 262-302.

Cologni, A. and M. Manera. 2008. Oil Prices, Inflation and Interest Rates in a Structural Cointegrated VAR Model for the G-7 Countries. *Energy Economics* 30: 856-888.

Cook, J. 2010. Don't Invert That Matrix. Accesed October 2, 2018. Retrieved from https://www.johndcook.com/blog/2010/01/19/dont-invert-that-matrix/

Cooney, K., J. Pater, and K. Meadows. 2008. Leading the Way: BPA's Efforts to Accelerate Energy Efficiency in the Northwest. Paper Presented at ACEEE Summer Study on Energy Efficiency in Buildings, Pacific Grove, August 17-22. Accessed October 11, 2016. Retrieved from http://www.aceee.org/files/proceedings/2008/data/papers/8_594.pdf

Cooper, J.P. and C.R. Nelson. 1975. The Ex Ante Prediction Performance of the St. Louis and FRB-MIT-PENN Econometric Models and Some Results on Composite Predictors. *Journal of Money, Credit, and Banking* 7: 1-32.

Cuaresma, J.C. and A. Breitenfellner. 2008. Crude Oil Prices and the Euro-Dollar Exchange Rate: A Forecasting Exercise. Working Paper, University of Innsbruck Working Papers in Economics and Statistics, Innsbruck. Accessed October 11, 2016. Retrieved from ftp://ftp.repec.org/opt/ReDIF/RePEc/inn/wpaper/2008-08.pdf

Dahl, C. and T.E. Duggan. 1998. Survey of Price Elasticities from Economic Exploration Models of US Oil and Gas Supply. *Journal of Energy Finance & Development* 3: 129-169.

Deaton, A. and J. Muellbauer. 1980. An Almost Ideal Demand System. *The American Economic Review* 70: 312-326.

Dèes, S., A. Gasteuil, R.K. Kaufmann, and M. Mann. 2008. Assessing the Factors Behind Oil Price Changes. Working paper, European Central Bank, Frankfurt. Accessed October 11, 2016. Retrieved from https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp855.pdf?83b595eb9f524cc671ce32c0341c81de

Del Negro, M. and G.E. Primiceri. 2015. Time Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum. *The Review of Economic Studies* 82: 1342-1345.

Demiralp, S. and K. Hoover. 2003. Searching for the Causal Structure of a Vector Autoregression. *Oxford Bulletin of Economics and Statistics* 65: 745-767.

Du, X., C.L. Yu, and D.J. Hayes. 2011. Speculation and Volatility Spillover in the Crude Oil and Agricultural Commodity Markets: A Bayesian Analysis. *Energy Economics* 33: 497-503.

Duangnate, K. 2015. Essays on the Dynamics of and Forecasting Ability within the US Energy Sector. Doctoral Dissertation, Texas A&M University, College Station.

Ederington, L.H., C.S. Fernando, T.K. Lee, S.C. Linn, and A.D. May. 2011. Factors Influencing Oil Prices: A Survey of the Current State of Knowledge in the Context of the 2007-08 Oil Price Volatility. Energy Information Administration, Washington, DC. Accessed October 11, 2016. Retrieved from http://www.eia.gov/finance/markets/reports_presentations/factors_influencing_oil_prices.pdf

Evergreen Economics. 2012. Northwest Ductless Heat Pump Initiative: Market Progress Evaluation Report #2. Accessed October 11, 2016. Retrieved from http://neea.org/docs/reports/northwest-ductless-heat-pump-initiative-market-progress-evaluation-report-2.pdf?sfvrsn=6

Federal Reserve Bank of St. Louis. 2018. *FRED Economic Data* [data files]. St Louis: Federal Reserve Bank of St. Louis. Accessed May 15, 2018. Retrieved from https://fred.stlouisfed.org/

Fernandez, V.P. 1999. Forecasting Home Appliance Sales: Incorporating Adoption and Replacement. *Journal of International Consumer Marketing* 12: 39-61.

Forni, M. and L. Gambetti. 2010. The Dynamic Effects of Monetary Policy: A Structural Factor Model Approach. *Journal of Monetary Economics* 57: 203-216.

Francisco, P.W., D. Baylon, B. Davis, and L. Palmiter. 2004. Heat Pump System Performance in Northern Climates. *American Society of Housing, Refrigerating and Air-Conditioning Engineers Transactions* 10: 442-451.

Free Software Foundation. 2018. *GNU Scientific Library* (version 2.1) [computer software]. Accessed November 19, 2018. Retrieved from https://www.gnu.org/software/gsl/

Friedman, N. 2016. How a Pipeline Leak in Alabama Can Move Gas Prices. *Wall Street Journal*, September 15, 2016. Accessed October 14, 2016. Retrieved from http://blogs.wsj.com/moneybeat/2016/09/15/how-a-pipeline-leak-in-alabama-can-move-gas-prices/

Gabel, H.L. 1979. A Simultaneous Equation Analysis of the Structure and Performance of the United States Petroleum Refining Industry. *The Journal of Industrial Economics* 28: 89-104.

Galarraga, I., L.M. Abadie, and A. Ansuategi. 2013. Efficiency Effectiveness and Implementaion Feasibility of Energy Efficiency Rebates: The "Renove" Plan in Spain. *Energy Economics* 40: S98-S107.

Gambetti, L. 2012. Lecture 2: Time-Varying Coefficients VARS. Universitat Autonoma de Barcelona, Barcelona. Accessed July 21, 2018. Retrieved from http://pareto.uab.es/lgambetti/Lecture2LBS.pdf

Gates, R.W. 1983. Investing in Energy Conservation: Are Homeowners Passing Up Yields? *Energy Policy* 11: 63-71.

Gelman, A., J.B. Carlin, H.S. Stern, D.B. Dunson, A. Vehtari, and D.B. Rubin. 2013. *Bayesian Data Analysis*. Boca Raton: CRC Press.

Ghatak, S., C. Milner, and U. Utkulu. 1997. Exports, Export Composition and Growth: Cointegration and Causality Evidence for Malaysia. *Applied Economics* 29: 213-223.

Gillingham, K., R.G. Newell, and K. Palmer. 2009. Energy Efficiency Economics and Policy. Working paper, National Bureau of Economic Research, Cambridge. Accessed October 11, 2016. Retrieved from http://www.nber.org/papers/w15031.

Golder, P.N., and G.J. Tellis. 2004. Growing, Growing, Gone: Cascades, Diffusion, and Turning Points in the Product Life Cycle. *Marketing Science* 23:207-218.

Guerra, S. 2008. Long Run Relationship between Oil Prices and Aggregate Oil Investment: Empirical Evidence. Working paper, United States Association for Energy Economics, Cleveland. Accessed October 11, 2016. Retrieved from http://www.usaee.org/usaee2007/submissions/OnlineProceedings/Sergio%20Guerra%20WP%2017-05-07.pdf

Günter, J.H.F. 2014. How Do Oil Producers Respond to Oil Demand Shocks? *Energy Economics* 44:1-13.

Hamilton, J.D. 1983. Oil and the Macroeconomy since World War II. *Journal of Political Economy* 91: 228-248.

Hamilton, J.D. 1994. *Time Series Analysis*. Princeton: Princeton University Press.

Hamilton, J.D. 2009. Understanding Crude Oil Prices. Working paper, National Bureau of Economic Research, Cambridge. Accessed July 8, 2018. Retrieved from http://www.nber.org/papers/w14492

Hamilton, J.D. and A.M. Herrera. 2004. Oil Shocks and Aggregate Macroeonomic Behavior: The Role of Monetary Policy. *Journal of Money, Credit, and Banking* 36: 17-37.

He, Y., S. Wang, and K.K. Lai. 2010. Global Economic Activity and Crude Oil Prices: A Cointegration Analysis. *Energy Economics* 32: 868-876.

Heeler, R.M. and T.P. Hustad. 1980. Problems in Predicting New Product Growth for Consumer Durables. *Management Science* 26:1007-1020.

Herrera, A.M. 2018. Oil Price Shocks, Inventories, ad Macroeconomic Dynamics. *Macroeconomic Dynamics* 22: 620:639.

Herrera, A.M, M.B. Karaki, and S.K. Rangaraju. 2017. Where Do Jobs Go When Oil Prices Drop? *Energy Economics* 64: 469-482.

Herrera, A.M. and E. Pesavento. 2009. Oil Price Shocks, Systematic Monetary Policy, and the 'Great Moderation.' *Macroeconomic Dynamics* 13: 107-137.

Hoff, P.D. 2009. *A First Course in Bayesian Statistical Methods*. Dordrecht: Springer.

Howarth, R.B. and B. Andersson. 1993. Market Barriers to Energy Efficiency. *Energy Economics* 15: 262-272.

Howarth, R.B. and A.H. Sanstad. 1995. Discount Rates and Energy Efficiency. *Contemporary Economic Policy* 13: 101-109.

Hsaio, C. 1979. Autoregressive Modeling of Canadian Money and Income Data. *Journal of the American Statistical Association* 74: 553-560.

Hsaio, C. 1981. Autoregressive Modeling and Money-Income Causality Detection. *Journal of Monetary Economics* 7: 85-106.

Huang, R.D., R.W. Masulis, and H.R. Stoll. 1996. Energy Shocks and Financial Markets. *Journal of Futures Markets* 16: 1-27.

Ipatova, E. 2014. Essays on Factor Models, Application to the Energy Markets. Doctoral Dissertation, City University of London, London.

Jaffe, A.B. and R.N. Stavins. 1994. The Energy-Efficiency Gap: What Does It Mean? *Energy Policy* 22: 804-810.

Jain, D.C. and R.C. Rao. 1990. Effect of Price on the Demand for Consumer Durables: Modeling, Estimation, and Findings. *Journal of Business & Economic Statistics* 8: 163-170.

Jones, C.M. and G. Kaul. 1996. Oil and the Stock Markets. *The Journal of Finance* 51: 463-491.

Juvenal, L. and I. Petrella. 2012. Speculation and the Oil Market. Working paper, Federal Reserve Bank of St. Louis, St. Louis. Accessed September 4, 2018. Retrieved from https://pdfs.semanticscholar.org/f65b/669c792956769947310f22ecce06fc4fa63e.pdf

Kalisch, M. and P. Buhlmann. 2007. Estimating High-Dimensional Directed Acyclic Graphs with the PC-Algorithm. *Journal of Machine Learning Research* 8: 613:636.

Kang, W., F. Perez de Gracia, and R.A. Ratti. Forthcoming. The Asymmetric Response of Gasoline Prices to Oil Price Shocks and Policy Uncertainty. *Energy Economics* (in press, corrected proof).

Karfakis, C.J. and D.M. Moschos. 1990. Interest Rate Linkages within the European Monetary System: A Time Series Analysis. *Journal of Money, Credit, & Banking* 22: 288-294.

Kaufmann, R.K., S. Dées, A. Gasteuil, and M. Mann. 2008. Oil Prices: The Role of Refinery Utilization, Futures Markets and Non-Linearities. *Energy Economics* 30: 2609-2622.

Kilian, L. 2009. Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review* 99: 1053-69.

Kilian, L. and H. Lutkepohl. 2017. *Structural Vector Autoregressive Analysis*. Acessed September 4, 2018. Retrieved from http://www-personal.umich.edu/~lkilian/book.html

Kilian, L. and D. Murphy. 2010. Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. Working paper, Department of Economics, University of Michigan, Ann Arbor. Accessed September 8, 2018. Retrieved from http://www-personal.umich.edu/~lkilian/km080210.pdf

Kilian, L. and R.J. Vigfusson. 2017. The Role of Oil Price Shocks in Causing US Recessions. *Journal of Money, Credit, and Banking* 49: 1747:1776.

Kim, S., N. Shephard, and S. Chib. 1998. Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Status* 65: 361-393.

Kim, C.J. and R. Nelson. 1999. *State-Space Models with Regime Switching.* Cambridge: The MIT Press.

Kitagawa, G. 1996. Monte Carlo Filter and Smoother for Gaussian Nonlinear State Space Models. *Journal of Computational and Graphical Statistics* 5: 1-25.

Korobilis, D. 2012. Assessing the Transmission of Monetary Policy Using Time-Varying Parameter Dynamic Factor Models. *Oxford Bulletin of Economics and Statistics* 75: 157-179.

Krichene, N. 2006. World Crude Oil Markets: Monetary Policy and the Recent Oil Shock. Working paper, International Monetary Fund, Washington, DC. Accessed October 11, 2016. Retrieved from https://www.imf.org/external/pubs/ft/wp/2006/wp0662.pdf

Kuper, G.H. 2012. Inventories and Upstream Gasoline Price Dynamics. *Energy Economics* 34: 208-214.

Levine, M.D., J.G. Koomey, J.E. McMahon, A.H. Sanstad, and E. Hirst. 1995. Energy Efficiency Policy and Market Failures. *Annual Review of Energy and the Environment* 20: 535-555.

Lippi, F. and A. Nobili. 2008. Oil and the Macroeconomy: A Structural VAR Analysis with Sign Restrictions. Working paper, Bank of Italy, Rome. Accessed September 4, 2018. Retrieved from https://www.bancaditalia.it/pubblicazioni/altri-atti-seminari/2008/Lippi_Nobili_8_04_08.pdf

Lippi, F. and A. Nobili. 2012. Oil and the Macroeconomy: A Quantitative Structural Analysis. *Journal of the European Economic Association* 10: 1059-1083.

Litterman, R. 1986. Specifying Vector Autoregressions for Macroeconomic Forecasting. *Federal Reserve Bank of Minneapolis Staff Report 92*. Accessed October 11,

2016. Retrieved from https://www.minneapolisfed.org/research/staff-reports/specifying-vector-autoregressions-for-macroeconomic-forecasting

Liu, X.M., H.Y. Song, and P. Romilly. 1997. An Empirical Investigation of the Causal Relationship between Openness and Economic Growth in China. *Applied Economics* 29: 1679-1686.

Lizardo, R.A. and A.V. Mollick. 2010. Oil Price Fluctuations and US Dollar Exchange Rates. *Energy Economics* 32: 399-408.

Loughran, D.S. and J. Kulick. 2004. Demand-Side Management and Energy Efficiency in the United States. *Energy Journal* 25:19-43.

Lund, P. 2006. Market Penetration Rates of New Energy Technologies. *Energy Policy* 34: 3317-3326.

Lutkepohl, H. 2005. *New Introduction to Multiple Time Series.* Berlin: Springer.

Lutz, C. 2015. The International Impact of US Unconventional Monetary Policy. *Applied Economics Letters* 22: 955-959.

Maslyuk, S. and R. Smyth. 2009. Non-Linear Unit Root Properties of Crude Oil Production. *Energy Economics* 31: 109-118.

McClure, T. 2016. How Kalman Filters Work, Part 2. Accessed October 2, 2018. Retrieved from http://www.anuncommonlab.com/articles/how-kalman-filters-work/part2.html

Melolinna, M. 2012. Macroeconomic Shocks in an Oil Market VAR. Working paper, European Central Bank, Frankfurt. Accessed September 16, 2018. Retrieved from https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp1432.pdf?24a4199c0d6dad77861710d3c6aae089

Michelsen, C.C. and R. Madlener. 2012. Homeowners' Preferences for Adopting Innovative Residential Heating Systems: A Discrete Choice Analysis for Germany. *Energy Economics* 34: 1271-1283.

Mills, B.F. and J. Schleich. 2010. Why Don't Households See the Light? Explaining the Diffusion of Compact Fluorescent Lamps. *Resource and Energy Economics* 32: 363-378.

Mills, E., and A. Rosenfeld. 1996. Consumer Non-Energy Benefits As a Motivation for Making Energy-Efficient Improvements. *Energy* 21: 707-720.

Mohn, K. 2008. Efforts and Efficiency in Oil Exploration: A Vector Error-Correction Approach. *The Energy Journal* 29: 53-78.

Mohn, K. and P. Osmundsen. 2008. Exploration Economics in a Regulated Petroleum Province: The Case of the Norwegian Continental Shelf. *Energy Economics* 30: 303-320.

Mu, X. 2007. Weather, Storage, and Natural Gas Price Dynamics: Fundamentals and Volatility. *Energy Economics* 29: 46-63.

Mumtaz, H., P. Zabczyk, and C. Ellis. 2011. What Lies Beneath? A Time-Varying FAVAR Model for the UK Transmission Mechanism. Working paper, European Central Bank, Frankfurt. Accessed October 11, 2016. Retrieved from https://www.ecb.europa.eu/pub/research/working-papers/html/papers-2011.en.html

Murray, A.G. and B.F. Mills. 2011. Read the Label! Energy Star Appliance Label Awareness and Uptake Among US Consumers. *Energy Economics* 33: 1103-1110.

National Bureau of Economic Research. 2018. US Business Cycle Expansions and Contractions. Accessed September 26, 2018. Retrieved from http://www.nber.org/cycles.html

National Resources Defense Council. 2010. Efficient Appliances Save Energy - and Money. Accessed July 15, 2014. Retrieved from http://www.nrdc.org/air/energy/fappl.asp

Nelson, C.R. 1972. The Prediction Performance of the FRB-MIT-PENN Model of the US Economy. *American Economic Review* 62: 902-917.

Northwest Ductless Heat Pump Project. 2012. Contractor Orientation. Northwest Energy Efficiency Alliance. Accessed July 15, 2014. Retrieved from http://energizecorvallis.org/wp-content/uploads/2013/03/Contractor-Orientation-1.16.13.pdf

Northwest Ductless Heat Pump Project 2014. Ductless Heating & Cooling Systems. Accessed October 11, 2016. Retrieved from http://www.goingductless.com

Northwest Energy Efficiency Alliance. 2013. Efficient Ductless Heat Pumps. Accessed October 11, 2016. Retrieved from http://neea.org/initiatives/residential/ductless-heat-pumps.

Olatubi, W.O. and S.C. No. 2003. On the Vulnerability of the Oil and Gas Industry to Oil Price Changes. *Atlantic Economic Journal* 31: 363-375.

Olsen, K. 2011. Price Discovery in the Natural Gas Markets of the United States and Canada. Master's thesis, Texas A&M University, College Station.

Olsen, K.K., J.W. Mjelde, and D.A. Bessler. 2015. Price Formulation and the Law of One Price in International Linked Markets: An Examination of the Natural Gas Markets in the USA and Canada. *The Annals of Regional Science* 54: 117-142.

OpenMP. 2018. The OpenMP API Specification for Parallel Programming [computer software]. Accessed November 19, 2018. Retrieved from https://www.openmp.org/

The OpenMPI Project. 2018. OpenMPI: Open Source High Performance Computing (verion 1.10.2) [computer software]. Accessed November 19, 2018. Retrieved from https://www.open-mpi.org/

Panzone. L.A. 2013. Saving Money vs. Investing Money: Do Energy Ratings Influence Consumer Demand for Energy Efficient Goods? *Energy Economics* 38: 51-63.

Pindyck, R.S. 2001. The Dynamics of Commodity Spot and Futures Markets: A Primer. *The Energy Journal* 22: 1-29.

Prest, B.C. 2018. Explanations for the 2014 Oil Price Decline: Supply or Demand? *Energy Economics* 74: 63-75.

Primiceri, G.E. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72:821-852.

Radchenko, S. 2005. Oil Price Volatility and the Asymmetric Response of Gasoline Prices to Oil Price Increases and Decreases. *Energy Economics* 27: 708-730.

Ratti, R.A. and J.L. Vespignani. 2015. Commodity Prices and BRIC and G3 Liquidity: A SFAVEC Approach. *Journal of Banking & Finance* 52: 18-33.

R Core Team. 2014. R: A Language and Environment for Statistical Computing (version 3.1) for Windows [computer software]. Accessed November 19, 2018. Retrieved from https://www.r-project.org/

R Core Team. 2018. R: A Language and Environment for Statistical Computing (version 3.4) for Windows and Linux [computer software]. Accessed November 19, 2018. Retrieved from https://www.r-project.org/

Rencher, A.C. and W.F. Christensen. 2012. *Methods of Multivariate Analysis*. Hoboken: John Wiley & Sons, Inc.

Ringlund, G.B., K.E. Rosendahl, and T. Skjerpen. 2008. Does Oilrig Activity React to Oil Price Changes? An Empirical Investigation. *Energy Economics* 30: 371-396.

Rogers, E.M. 1962. *Diffusion of Innovations*. New York: The Free Press.

Rubio-Ramirez, J.F., D.F. Waggoner, and T. Zha. 2010. Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. *Review of Economic Studies* 77: 665-696.

Sadorsky, P. 1999. Oil Price Shocks and Stock Market Activity. *Energy Economics* 21: 449-469.

Şahin, A.Ş., B. Kılıç, and U. Kılıç. 2011. Optimization of Heat Pump Using Fuzzy Logic and Genetric Algorithm. *Heat and Mass Transfer* 47: 1553-1560.

Saito, M. and M. Matsumoto. 2013. Double-Precision SIMD-oriented Fast Mersenne Twister (SFMT) version 2.2.3. Accessed November 19, 2018. Retrieved from http://www.math.sci.hiroshima-u.ac.jp/~m-mat/MT/SFMT/index.html#dSFMT

SAS Institute Inc. 2013. SAS version 9.3 for Windows. Accessed November 19, 2018. Retrieved from https://www.sas.com/en_us/home.html

Schmittlein, D.C. and V. Mahajan. 1982. Maximum Likelihood Estimation for an Innovation Diffusion Model of New Product Acceptance. *Marketing Science* 1:57-78.

Srinivasan, V. and C.H. Mason. 1986. Nonlinear Least Squares Estimation of New Product Diffusion Models. *Marketing Science* 5:169-178.

Sims, C. 1980. Macroeconomics and Reality. *Econometrica* 48: 1-48.

Smith, J.L. and T.K. Lee. 2017. The Price Elasticity of U.S. Shale Oil Reserves. Energy Economics 67: 121-135.

Spiegelhalter, D.J., N.G. Best, B.P. Carlin, and A. van der Linde. 2002. Bayesian Measures of Model Complexity and Fit. *Journal of the Royal Statistical Society B* 64: 583-639.

Spirtes, P., C. Glymour, and R. Scheines. 2000. *Causation, Prediction, and Search*. Cambridge: The MIT Press.

Stecher, D. and K. Allison. 2012. Maximum Residential Energy Efficiency: Performance Results from Long-Term Monitoring of a Passive House. *American Society of Housing, Refrigerating and Air-Conditioning Engineers Transactions* 118: 127-134.

Stock, J. and M. Watson. 2002a. Forecasting Using Principal Components from a Large Number of Predictors. *Journal of the American Statistical Association* 97: 1167-1179.

Stock, J. and M. Watson. 2002b. Has the Business Cycle Changed and Why? Working paper, National Bureau of Economic Research, Cambridge. Accessed October 11, 2016. Retrieved from http://www.nber.org/papers/w9127

Stock, J. and M. Watson. 2002c. Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business Economics and Statistics* 20: 147-162.

Storm, P., D. Baylon, A. Armstrong, and J. Harris. 2012. Integrated Ductless Heat Pump Analysis: Developing an Emerging Technology into a Regional Efficiency Resource. Paper presentated at ACEEE Summer Study on Energy Efficiency in Buildings, Pacific Grove, August 12-17. Accessed October 11, 2016. Retrieved from http://aceee.org/files/proceedings/2012/start.htm

Sultan, F., J.U. Farley, and D.R. Lehmann. 1990. A Meta-Analysis of Diffusion Models. *Journal of Marketing Research* (27): 70–77.

Sutherland, J. 2012. Ductless Heat Pumps. *Journal of Light Construction* 30: 10.

Swift, J.R. and R.A. Meyer. 2010. Ductless Heat Pumps for Residential Customers in Connecticut. Paper presented at ACEEE Summer Study on Energy Efficiency in Buildings, Pacific Grove, August 15-20. Accessed October 11, 2016. Retrieved from http://aceee.org/files/proceedings/2010/data/papers/1960.pdf

Tan, H.B., E.T. Cheah, J.E.V. Johnson, M.C. Sung, and C.H. Chuah. 2012. Stock Market Capitalization and Financial Integration in the Asia Pacific Region. *Applied Economics* 44: 1951-1961.

Taylor, C.R. 1993. Dynamic Programming and the Curses of Dimensionality In C.R. Taylor, ed. *Applications of Dynamic Programming to Agricultural Decision Problems.* Boulder: Westview Press. 1-10.

The Tetrad Project. 2017. Tetrad: Graphical Causal Models. Accessed November 9, 2018. Retrieved from http://www.phil.cmu.edu/tetrad/

Tietenberg, T. and L. Lewis. 2012. Energy: The Transition from Depletable to Renewable Resources. Chapter 7 in *Environmental & Natural Resource Economics* 9th ed. Boston, Pearson, 140-179.

Tsamardinos, I., L.I. Brown, and C.F. Aliferis. 2006. The Max-Min Hill-Climbing Bayesian Network Structure Learning Algorithm. *Journal of Economic Perspectives* 65: 31-78.

Uhlig, H. 2005. What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics* 52: 381-419.

Uhlig, H. 2017. Shocks, Signs, Restrictions, and Identification. In B. Honore, A. Pakes, M. Pizzesi, and L. Samuelson, eds. *Advances in Economics and Econometrics Vol 1 and 2: Eleventh World Congress*. Cambridge: Cambridge University Press, 95-127.

Unalmis, D., I. Unalmis, and D.F. Unsal. 2012. On Oil Price Shocks: The Role of Storage. *IMF Economic Review* 60: 505-532.

United Energy Information Administration. 2018. *Petroleum & Other Liquids* [data files]. Accessed November 19, 2018. Retrieved from https://www.eia.gov/petroleum/

Vahid, F. and J.V. Issler. 2002. The Importance of Common Cyclical Features in VAR Analysis: a Monte-Carlo Study. *Journal of Econometrics* 109: 341-363.

Van den Bulte, C. and S. Stremersch. 2004. Social Contagion and Income Heterogeneity in New Production Diffusion: A Meta-Analytic Test. *Marketing Science* 23(4): 530-544.

Wang, Q. and X. Sun. 2017. Crude Oil Price: Demand, Supply, Economic Activity, Economic Policy Uncertainty and Wars – From the Perspective of Structure Equation Modelling (SEM). *Energy* 133: 483-490.

Wasi, N. and R.T. Carson. 2013. The Influence of Rebate Programs on the Demand for Water Heaters: The Case of New South Wales. *Energy Economics* 40: 645-656.

Working, H. 1949. The Theory of the Price of Storage. *American Economic Review* 39: 1254-1262.

Xianyi, Z. 2018. OpenBLAS: An Optimized BLAS Library (version 0.2.15) [computer software]. Accessed November 19, 2018. Retrieved from https://www.openblas.net/

Ye, M., J. Zyren, and J. Shore. 2002. Forecasting Crude Oil Spot Price Using OECD Petroleum Inventory Levels. *International Advances in Economic Research* 8: 324-333.

Zagalia, P. 2010. Macroeconomic Factors and Oil Futures Prices: A Data-Rich Model. *Energy Economics* 32: 409-417.

APPENDIX A

TABLES AND FIGURES

Table 2.1 Summary Statistics for Ductless Heat Pump Installations by Month

	Number of Installations		Total Per-Unit Cost		NEEA Expenditures	
Month	Households	Std. Dev.	Dollars	Std. Dev.	Dollars	Std. Dev.
January	169	89	4,206	207	135,073	36,335
February	148	40	4,371	322	130,167	30,489
March	214	70	4,323	482	191,446	57,157
April	174	22	4,563	563	129,412	52,945
May	201	54	4,611	564	128,821	54,135
June	213	57	4,562	386	153,766	36,191
July	282	89	4,374	360	117,287	42,220
August	299	132	4,640	664	147,338	36,442
September	373	171	4,284	214	142,923	33,510
October	352	155	4,359	350	144,919	26,696
November	403	179	4,413	444	163,220	41,383
December	648	261	4,465	366	354,877	40,818

The average value for each variable is shown by month; the standard deviations are given by Std. Dev. The data cover January 2009 through August 2013. All dollar amounts are CPI-adjusted. Household includes single-family households heated by electricity. Total per-unit cost represents all costs of a DHP installation, including but not limited to, equipment, labor, electrical, permit, tax, etc. This cost is before any rebates and/or tax credits.

Table 2.2 Parameter Estimates for the Adoption Model

Parameter	Coefficient Estimate	p-value
Coefficient of Innovation (α)	0.002	0.176
Coefficient of Imitation (β)	0.068	0.001
Market Potential (M)	329,442	0.395
Net Installation Costs	-1.495	0.042
NEEA Expenditures	0.316	0.223
Tax Code Change (= 1 if before the change)	0.930	0.085
Constant ($\hat{c}$)	4.783	0.322
February	-0.218	0.440
March	0.102	0.710
April	0.019	0.958
May	0.209	0.560
June	0.192	0.582
July	0.487	0.161
August	0.545	0.127
September	0.691	0.067
October	0.591	0.081
November	0.751	0.033
December	1.185	0.009
AR(1)	0.894	0.000

Adoption model estimated using nonlinear least squares in SAS procedure PROC MODEL. The convergence criterion is set to 0.0000001. The estimation procedure uses 55 observations. Root mean squared error is 61.133 and the mean absolute percentage error is 19 percent.

Table 2.3 Assumptions for Five-Year Forecast Simulations

Parameter	Value
Cost per Outdoor Unit	
Mean	$4,435
Standard Deviation	$414
NEEA Annual Expenditures	
2013	$0
2014	$1,857,000
2015	$1,116,000
2016	$1,396,000
2017	$1,856,000
2018	$1,207,000
Monthly Share of NEEA's Annual Expenditures	
January	0.072
February	0.071
March	0.088
April	0.076
May	0.059
June	0.080
July	0.053
August	0.071
September	0.076
October	0.077
November	0.086
December	0.193
Average Tax Credit	$500
NWDHPP Installation Rebate	$1,500
NWDHPP Annual Rebate Budgets	
2013	$3,903,541
2014	$5,289,868
2015	$5,103,705
2016	$4,924,095
2017	$4,750,805
2018	$17,390,363
Beginning Cumulative Installations	15,662
Maximum Market Potential	1,658,148

Cost per outdoor unit is taken as the median value of all installations in August 2013. Budgets for NEEA expenditures and rebates provided by NEEA. All dollar amounts are assumed to be in real dollars. Monthly shares of NEEA's annual expenditures calculated from NEEA's 2010 through 2012 expenditure data. Maximum market potential is an estimate of the number of electrically-heated single-family houses in the four-state region.

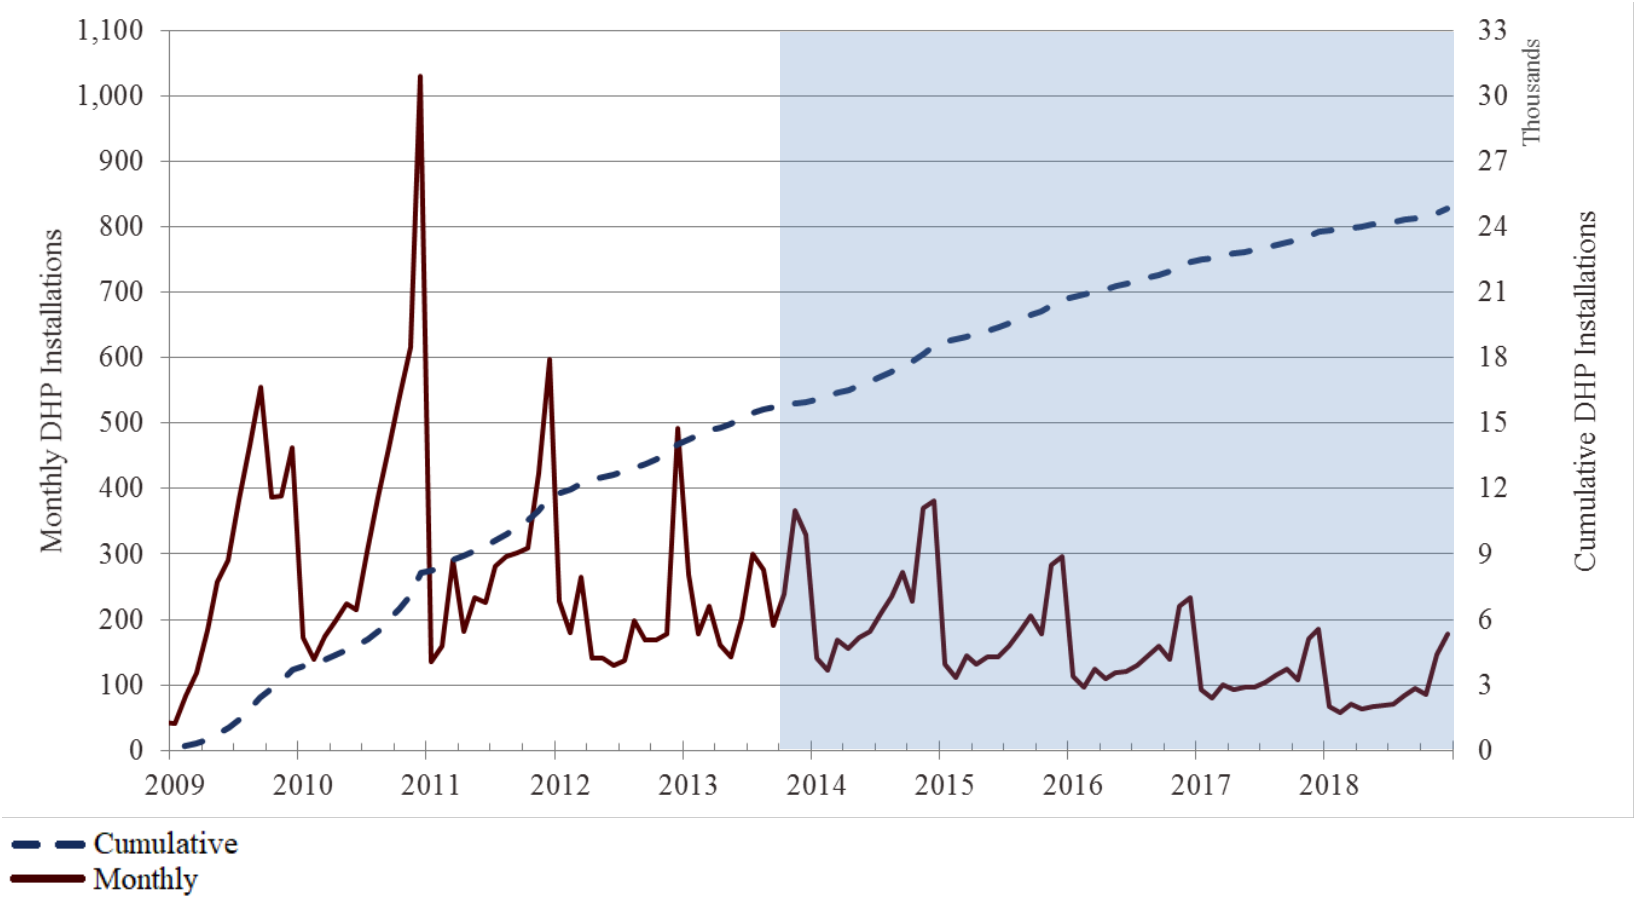

Figure 2.1 Monthly and cumulative installations of ductless heat pumps, 2009 through August 2013, with installations forecasted through 2018 (shaded area) using proposed annual rebate and marketing budgets

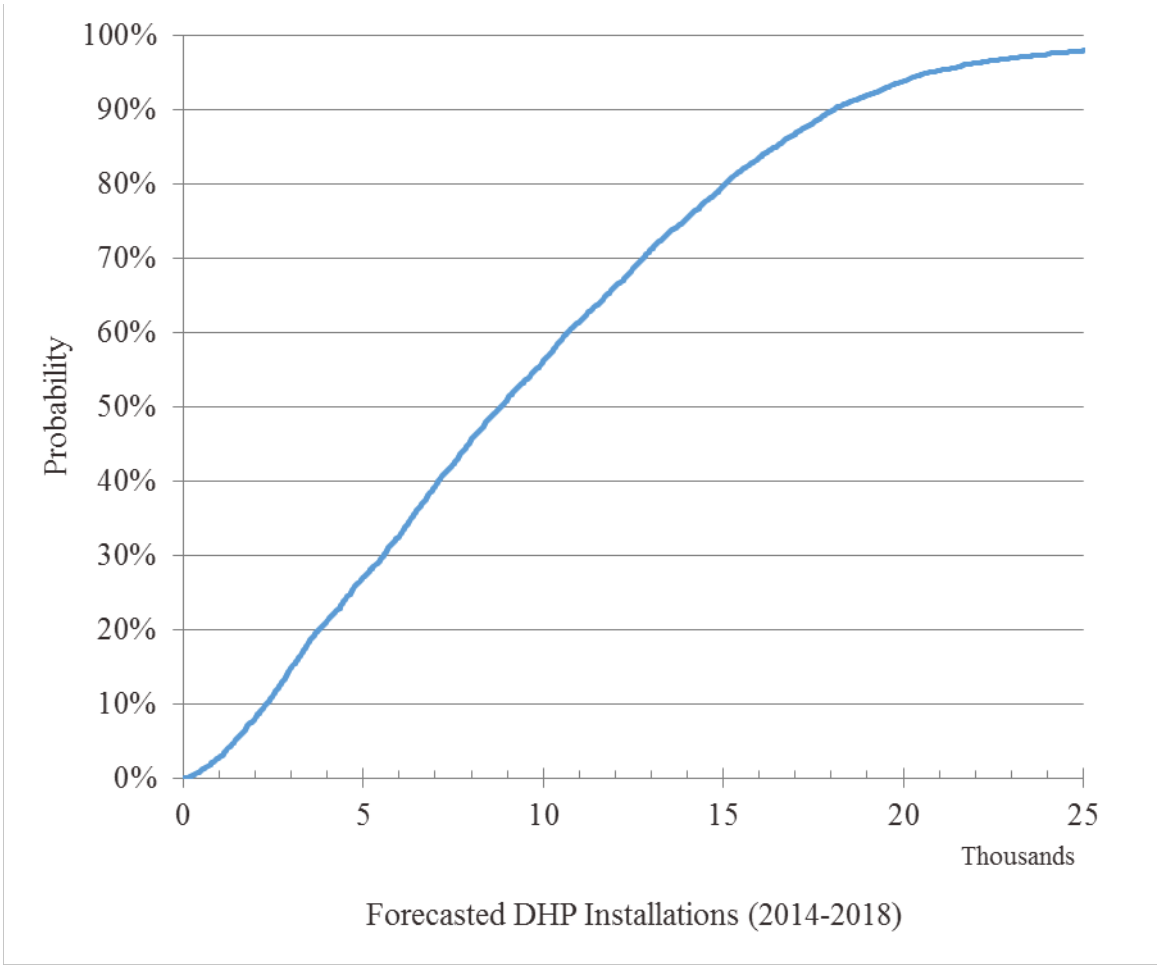

Figure 2.2 Cumulative distribution function for forecasted additional ductless heat pump installations through 2018 assuming base scenario

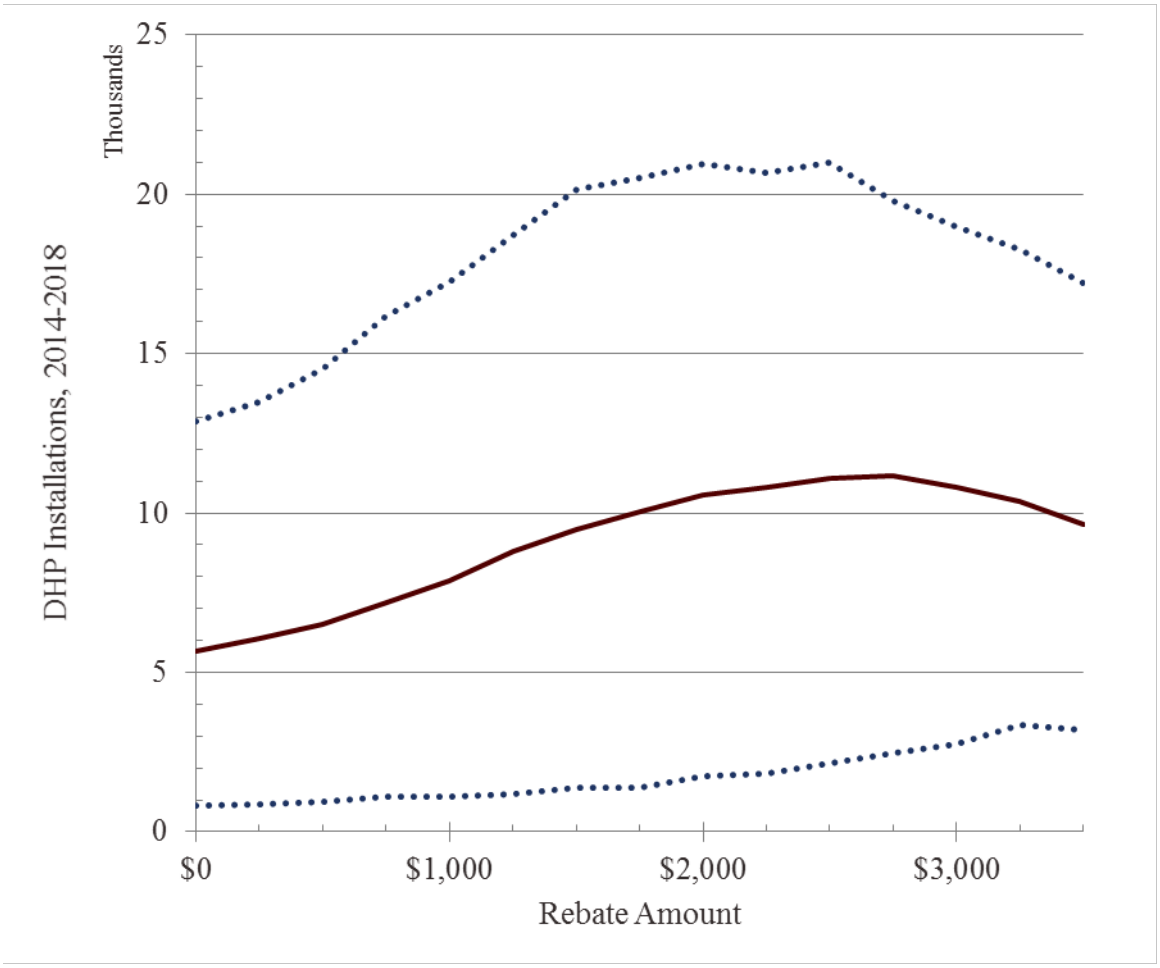

Figure 2.3 Expected and 95-percent confidence interval for forecasted additional ductless heat pump installations through 2018 for different rebate amounts

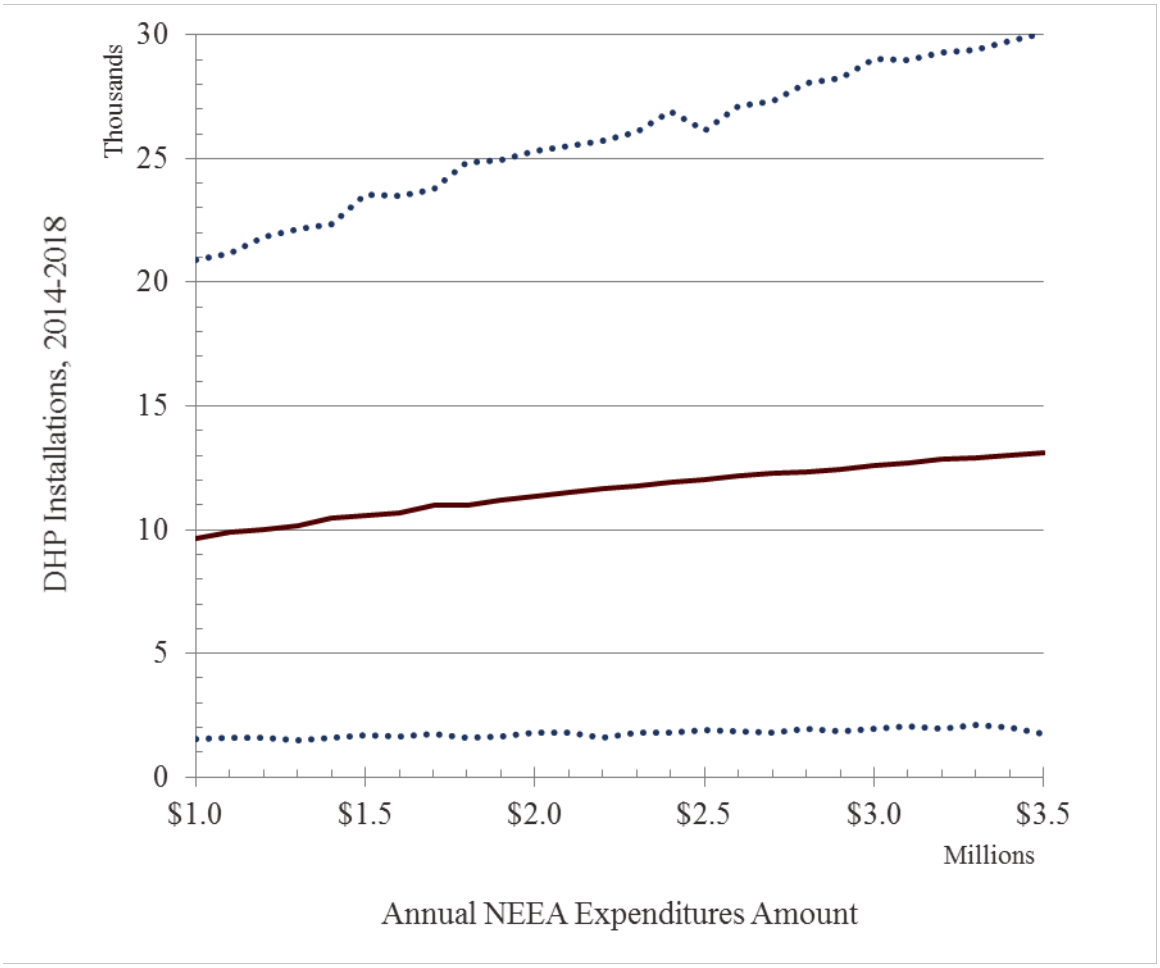

Figure 2.4 Expected and 95-percent confidence interval for forecasted additional ductless heat pump installations through 2018 for different annual NEEA expenditures on marketing and installer training

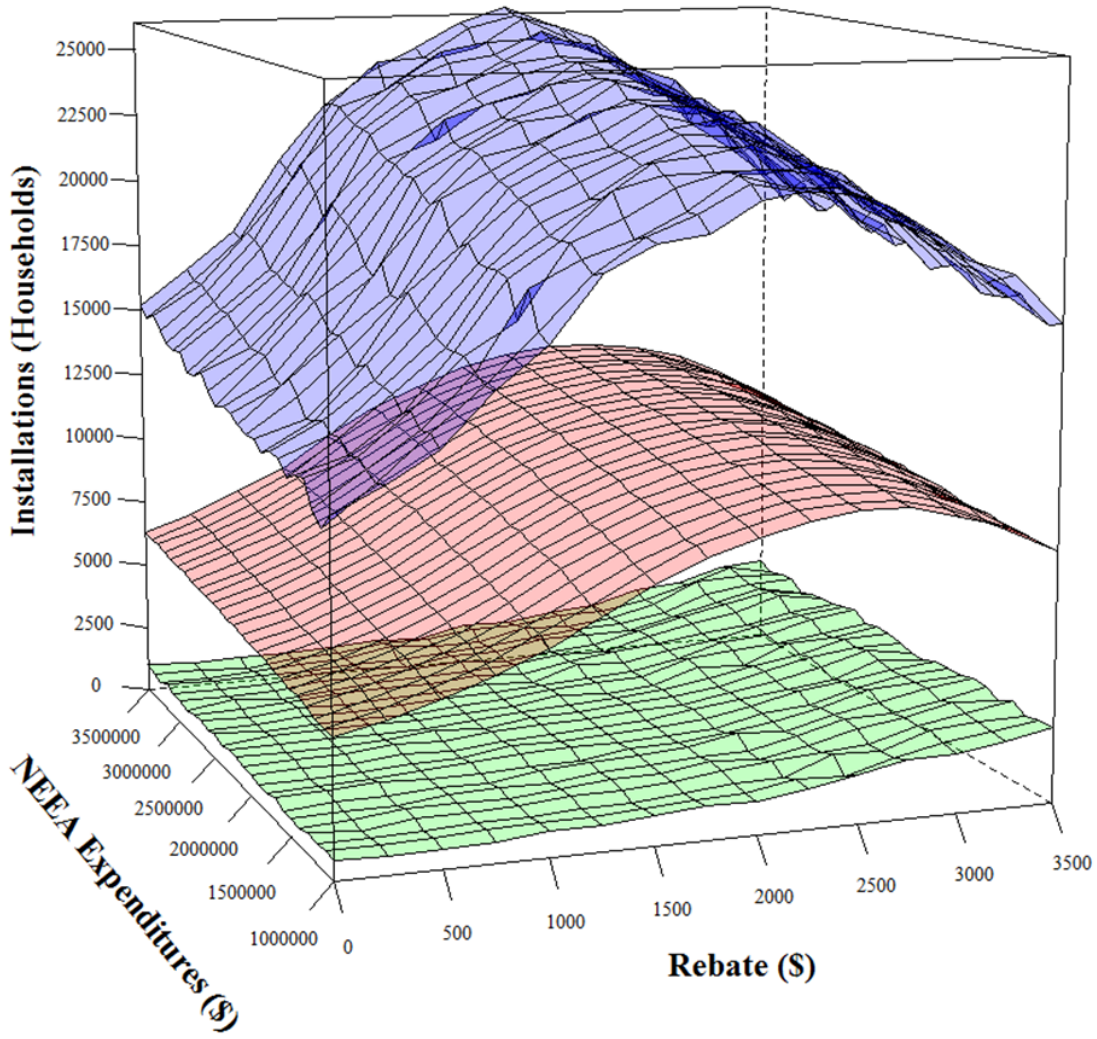

Figure 2.5 Expected and 95-percent confidence interval for additional installations through 2018 for different rebate amounts and levels of annual NEEA expenditures

Table 3.1 Data Series, Transformations, and Sources

Transformations: 1) no transformation; 2) first difference; 3) log; 4) first difference of log-transformation
Sources: EIA, US Energy Information Administration; FRED, US Federal Reserve Bank of St. Louis
** denotes observational series in the estimated FAVAR model*

Series	Trans.	Source
Industrial Production Index	4	FRED
Disposable Personal Income	4	FRED
Manufacturers Sales	4	FRED
Retailers Sales	4	FRED
US Crude Oil Rotary Rigs in Operation	4	EIA
US Percent Utilization of Refinery Operable Capacity	1	EIA
US Field Production of Crude Oil (Thousand Barrels)	4	EIA
US Imports from Persian Gulf Countries of Crude Oil (Thousand Barrels)	4	EIA
US Imports from OPEC Countries of Crude Oil (Thousand Barrels)	4	EIA
US Imports from Non-OPEC Countries of Crude Oil (Thousand Barrels)	4	EIA
US Imports from OPEC Countries of Total Petroleum Products (Thousand Barrels)	4	EIA
US Imports from Non-OPEC Countries of Total Petroleum Products (Thousand Barrels)	4	EIA
US Exports of Crude Oil (Thousand Barrels)	4	EIA
US Exports of Finished Petroleum Products (Thousand Barrels)	4	EIA
US Ending Stocks of Crude Oil in SPR (Thousand Barrels)	4	EIA
US Crude Oil Stocks at Tank Farms and Pipelines (Thousand Barrels)	4	EIA
US Crude Oil Stocks in Transit (on Ships) from Alaska (Thousand Barrels)	4	EIA
US Crude Oil Stocks at Refineries (Thousand Barrels)	4	EIA
US Product Supplied of Crude Oil and Petroleum Products (Thousand Barrels)	4	EIA
US Product Supplied of Hydrocarbon Gas Liquids (Thousand Barrels)	4	EIA
US Product Supplied of Finished Petroleum Products (Thousand Barrels)	4	EIA
US Ending Stocks of Total Petroleum Products (Thousand Barrels)	4	EIA
US Ending Stocks of Hydrocarbon Gas Liquids (Thousand Barrels)	4	EIA
US Ending Stocks of Unfinished Oils (Thousand Barrels)	4	EIA

Table 3.1 Continued

Series	Trans.	Source
US Ending Stocks of Gasoline Blending Components (Thousand Barrels)	4	EIA
US Ending Stocks of Finished Motor Gasoline (Thousand Barrels)	4	EIA
US Ending Stocks of Conventional Motor Gasoline (Thousand Barrels)	4	EIA
US Ending Stocks of Aviation Gasoline (Thousand Barrels)	4	EIA
US Ending Stocks of Kerosene-Type Jet Fuel (Thousand Barrels)	4	EIA
US Ending Stocks of Kerosene (Thousand Barrels)	4	EIA
US Ending Stocks of Distillate Fuel Oil (Thousand Barrels)	4	EIA
US Ending Stocks of Residual Fuel Oil (Thousand Barrels)	4	EIA
US Ending Stocks of Lubricants (Thousand Barrels)	4	EIA
US Ending Stocks of Waxes (Thousand Barrels)	4	EIA
US Ending Stocks of Petroleum Coke (Thousand Barrels)	4	EIA
US Ending Stocks of Asphalt and Road Oil (Thousand Barrels)	4	EIA
US Ending Stocks of Miscellaneous Petroleum Products (Thousand Barrels)	4	EIA
US Total Gasoline Retail Sales by Refiners (Thousand Gallons per Day)	4	EIA
New York Harbor Conventional Gasoline Regular Spot Price FOB (Dollars per Gallon)	4	EIA
US Gulf Coast Conventional Gasoline Regular Spot Price FOB (Dollars per Gallon)	4	EIA
New York Harbor No. 2 Heating Oil Spot Price FOB (Dollars per Gallon)	4	EIA
US Gulf Coast Kerosene-Type Jet Fuel Spot Price FOB (Dollars per Gallon)	4	EIA
Mont Belvieu, TX Propane Spot Price FOB (Dollars per Gallon)	4	EIA
Cushing, OK Crude Oil Future Contract 1 (Dollars per Barrel)	4	EIA
Cushing, OK Crude Oil Future Contract 2 (Dollars per Barrel)	4	EIA
Cushing, OK Crude Oil Future Contract 3 (Dollars per Barrel)	4	EIA
Cushing, OK Crude Oil Future Contract 4 (Dollars per Barrel)	4	EIA
New York Harbor No. 2 Heating Oil Future Contract 1 (Dollars per Gallon)	4	EIA
New York Harbor No. 2 Heating Oil Future Contract 3 (Dollars per Gallon)	4	EIA
All Employees: Manufacturing	4	FRED
All Employees: Construction	4	FRED

Table 3.1 Continued

Series	Trans.	Source
All Employees: Government	4	FRED
All Employees: Retail Trade	4	FRED
All Employees: Education and Health Services	4	FRED
All Employees: Financial Activities	4	FRED
All Employees: Professional and Business Services	4	FRED
All Employees: Durable Goods	4	FRED
All Employees: Mining and Logging	4	FRED
All Employees: Leisure and Hospitality	4	FRED
All Employees: Wholesale Trade	4	FRED
All Employees: Private Service-Providing	4	FRED
All Employees: Nondurable Goods	4	FRED
All Employees: Utilities	4	FRED
All Employees: Other Services	4	FRED
All Employees: Information	4	FRED
All Employees: Mining and Logging: Oil and Gas Extraction	4	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Construction	4	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Manufacturing	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Goods-Producing	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Goods-Producing	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Private Service-Providing	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Private Service-Providing	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Mining and Logging	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Mining and Logging	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Construction	1	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Manufacturing	1	FRED
Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Manufacturing	1	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Durable Goods	1	FRED

Table 3.1 Continued

Series	Trans.	Source
Average Hourly Earnings of Production and Nonsupervisory Employees: Durable Goods	4	FRED
Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Durable Goods	1	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Nondurable Goods	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Nondurable Goods	4	FRED
Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Nondurable Goods	1	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Trade, Transportation, and Utilities	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Trade, Transportation, and Utilities	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Wholesale Trade	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Wholesale Trade	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Retail Trade	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Retail Trade	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Transportation and Warehousing	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Transportation and Warehousing	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Utilities	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Utilities	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Information	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Information	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Financial Activities	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Financial Activities	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Professional and Business Services	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Professional and Business Services	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Education and Health Services	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Education and Health Services	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Leisure and Hospitality	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Leisure and Hospitality	4	FRED
Average Weekly Hours of Production and Nonsupervisory Employees: Other	1	FRED
Average Hourly Earnings of Production and Nonsupervisory Employees: Other Services	4	FRED

Table 3.1 Continued

Series	Trans.	Source
Number of Civilians Unemployed for Less Than 5 Weeks	1	FRED
Number of Civilians Unemployed for 27 Weeks and Over	1	FRED
Number of Civilians Unemployed for 5 to 14 Weeks	1	FRED
Number of Civilians Unemployed for 15 to 26 Weeks	1	FRED
Civilian Unemployment Rate	1	FRED
Consumer Price Index for All Urban Consumers: All Items	4	FRED
Consumer Price Index for All Urban Consumers: Apparel	4	FRED
Consumer Price Index for All Urban Consumers: Food and Beverages	4	FRED
Consumer Price Index for All Urban Consumers: Housing	4	FRED
Consumer Price Index for All Urban Consumers: Medical Care	4	FRED
Consumer Price Index for All Urban Consumers: Transportation	4	FRED
Consumer Price Index for All Urban Consumers: Durables	4	FRED
Consumer Price Index for All Urban Consumers: Transportation services	4	FRED
Consumer Price Index for All Urban Consumers: Tuition, other school	4	FRED
Consumer Price Index for All Urban Consumers: Sugar and sweets	4	FRED
Consumer Price Index for All Urban Consumers: Tobacco and smoking	4	FRED
Consumer Price Index for All Urban Consumers: Electricity	4	FRED
Consumer Price Index for All Urban Consumers: Professional services	4	FRED
Consumer Price Index for All Urban Consumers: New vehicles	4	FRED
Consumer Price Index for All Urban Consumers: Airline fare	4	FRED
Chicago Fed National Activity Index: Sales, Orders and Inventories	2	FRED
Capacity Utilization: Total Industry	2	FRED
Vehicle Miles Traveled	4	FRED
Real personal consumption expenditures: Durable goods	4	FRED
Real personal consumption expenditures: Food	4	FRED
Real personal consumption expenditures: Nondurable goods	4	FRED
Real personal consumption expenditures: Energy goods and services	4	FRED

Table 3.1 Continued

Series	Trans.	Source
Real personal consumption expenditures: Services	4	FRED
Privately Owned Housing Starts: 1-Unit Structures	4	FRED
Housing Starts: 2-4 Units	3	FRED
Privately Owned Housing Starts: 5-Unit Structures or More	4	FRED
Housing Starts in Midwest Census Region	4	FRED
Housing Starts in Northeast Census Region	4	FRED
Housing Starts in South Census Region	4	FRED
Housing Starts in West Census Region	4	FRED
Industrial Production: Final products	4	FRED
Industrial Production: Consumer goods	4	FRED
Industrial Production: Durable consumer goods	4	FRED
Industrial Production: Miscellaneous durable goods	4	FRED
Industrial Production: Miscellaneous nondurable goods	4	FRED
Industrial Production: Consumer energy products	4	FRED
Industrial Production: Residential utilities	4	FRED
Industrial Production: Industrial and other equipment	4	FRED
Industrial Production: Industrial equipment	4	FRED
Industrial Production: Other equipment	4	FRED
Industrial Production: Defense and space equipment	4	FRED
Industrial Production: Consumer parts	4	FRED
Industrial Production: Equipment parts	4	FRED
Industrial Production: Computer and other board assemblies and parts	4	FRED
Industrial Production: Other equipment parts	4	FRED
Industrial Production: Other nondurable materials	4	FRED
Industrial Production: Miscellaneous nondurable materials	4	FRED
Industrial Production: Nonindustrial supplies	4	FRED
Industrial Production: Construction supplies	4	FRED

Table 3.1 Continued

Series	Trans.	Source
Industrial Production: Business supplies	4	FRED
Industrial Production: General business supplies	4	FRED
Industrial Production: Finished processing	4	FRED
Industrial Production: Final Products and Nonindustrial Supplies	4	FRED
Industrial Production: Mining	4	FRED
Industrial Production: Electric and gas utilities	4	FRED
Industrial Production: Durable manufacturing	4	FRED
Industrial Production: Other manufacturing	4	FRED
Industrial Production: Computers, communications equipment, and	4	FRED
Industrial Production: Manufacturing (SIC)	4	FRED
Producer Price Index for All Commodities	4	FRED
Producer Price Index by Industry: Total Manufacturing Industries	4	FRED
Producer Price Index by Industry: Total Mining Industries	4	FRED
Producer Price Index by Industry: Oil and Gas Field Machinery and Equipment Manufacturing	4	FRED
Producer Price Index by Industry: Support Activities for Oil and Gas Operations	4	FRED
Producer Price Index by Industry: Drilling Oil and Gas Wells	4	FRED
Effective Federal Funds Rate*	1	FRED
3-Month Treasury Bill: Secondary Market Rate	1	FRED
6-Month Treasury Bill: Secondary Market Rate	1	FRED
Monetary Base; Total	4	FRED
M1 Money Stock	4	FRED
M2 Money Stock	4	FRED
Total Reserves of Depository Institutions	4	FRED
Canada / US Foreign Exchange Rate	4	FRED
China / US Foreign Exchange Rate	4	FRED
Japan / US Foreign Exchange Rate	4	FRED
Switzerland / US Foreign Exchange Rate	4	FRED

Table 3.1 Continued

Series	Trans.	Source
US / Australia Foreign Exchange Rate	4	FRED
US / U.K. Foreign Exchange Rate	4	FRED
1-Year Treasury Constant Maturity Rate	1	FRED
10-Year Treasury Constant Maturity Rate	1	FRED
5-Year Treasury Constant Maturity Rate	1	FRED
Commercial and Industrial Loans, All Commercial Banks	4	FRED
Consumer Loans, All Commercial Banks	4	FRED
Cushing, OK WTI Spot Price FOB (Dollars per Barrel)*	4	FRED

Table 3.2 Values of the Deviance Information Criterion (DIC)

Latent Factors	Lags in Transition Equation	DIC
1	2	230,357
	3	219,797
	4	252,678
	6	261,513
	12	234,266
2	2	265,708
	3	350,914
	4	330,155
	6	302,747
3	2	6,128,302
	3	423,844
	4	425,358
4	2	25,869,065
	3	3,368,701

All DIC values are calculated from log-likelihoods estimated by particle filter using 1,000 particles.

Table 3.3 Median R-Squared Values for Ordinary Least Squares-Estimated Regressions of Each Informational Series (Individually) on the Latent Factor

Category	Transformed Series Name	R-Squared
Activity Measures	Percentage Change in Manufacturers Sales	0.379
	Percentage Change in Retailers Sales	0.222
	Percentage Change in Vehicle Miles Traveled	0.109
	Change in Capacity Utilization: Total Industry	0.060
	Percentage Change in Real personal consumption expenditures: Services	0.031
	Percentage Change in Real personal consumption expenditures: Durable goods	0.019
	Percentage Change in Real personal consumption expenditures: Nondurable goods	0.016
	Change in Chicago Fed National Activity Index: Sales, Orders and Inventories	0.012
	Percentage Change in Real personal consumption expenditures: Food	0.006
	Percentage Change in Disposable Personal Income	0.003
	Percentage Change in Real personal consumption expenditures: Energy goods and services	0.000
Consumer Price Indices	Percentage Change in Consumer Price Index for All Urban Consumers: Electricity	0.070
	Percentage Change in Consumer Price Index for All Urban Consumers: Sugar and sweets	0.068
	Percentage Change in Consumer Price Index for All Urban Consumers: New vehicles	0.053
	Percentage Change in Consumer Price Index for All Urban Consumers: Transportation	0.040
	Percentage Change in Consumer Price Index for All Urban Consumers: Apparel	0.033
	Percentage Change in Consumer Price Index for All Urban Consumers: Medical Care	0.029
	Percentage Change in Consumer Price Index for All Urban Consumers: Food and Beverages	0.027
	Percentage Change in Consumer Price Index for All Urban Consumers: Airline fare	0.025
	Percentage Change in Consumer Price Index for All Urban Consumers: Tobacco and smoking	0.020
	Percentage Change in Consumer Price Index for All Urban Consumers: All Items	0.014
	Percentage Change in Consumer Price Index for All Urban Consumers: Tuition, other school	0.013
	Percentage Change in Consumer Price Index for All Urban Consumers: Housing	0.008
	Percentage Change in Consumer Price Index for All Urban Consumers: Professional services	0.008
	Percentage Change in Consumer Price Index for All Urban Consumers: Durables	0.002

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Consumer Price Indices	Percentage Change in Consumer Price Index for All Urban Consumers: Transportation services	0.000
Employment: Earnings	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Private Service-Providing	0.172
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Trade, Transportation, and Utilities	0.156
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Nondurable Goods	0.113
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Retail Trade	0.112
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Professional and Business Services	0.090
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Wholesale Trade	0.071
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Education and Health Services	0.057
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Transportation and Warehousing	0.039
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Goods-Producing	0.032
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Financial Activities	0.029
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Mining and Logging	0.024
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Utilities	0.022
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Information	0.021
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Construction	0.018
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Durable Goods	0.014
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Other Services	0.004

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Employment: Earnings	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Manufacturing	0.001
	Percentage Change in Average Hourly Earnings of Production and Nonsupervisory Employees: Leisure and Hospitality	0.001
Employment: Employees by Industry	Percentage Change in All Employees: Manufacturing	0.637
	Percentage Change in All Employees: Durable Goods	0.569
	Percentage Change in All Employees: Professional and Business Services	0.542
	Percentage Change in All Employees: Wholesale Trade	0.539
	Percentage Change in All Employees: Nondurable Goods	0.527
	Percentage Change in All Employees: Private Service-Providing	0.479
	Percentage Change in All Employees: Construction	0.378
	Percentage Change in All Employees: Financial Activities	0.305
	Percentage Change in All Employees: Leisure and Hospitality	0.290
	Percentage Change in All Employees: Mining and Logging	0.285
	Percentage Change in All Employees: Information	0.271
	Percentage Change in All Employees: Other Services	0.268
	Percentage Change in All Employees: Retail Trade	0.213
	Percentage Change in All Employees: Utilities	0.054
	Percentage Change in All Employees: Mining and Logging: Oil and Gas Extraction	0.041
	Percentage Change in All Employees: Government	0.040
	Percentage Change in All Employees: Education and Health Services	0.013
Employment: Hours	Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Durable Goods	0.004
	Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Manufacturing	0.002
	Average Weekly Overtime Hours of Production and Nonsupervisory Employees: Nondurable Goods	0.001
	Average Weekly Hours of Production and Nonsupervisory Employees: Financial Activities	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Education and Health Services	0.000

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Employment: Hours	Average Weekly Hours of Production and Nonsupervisory Employees: Professional and Business Services	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Information	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Utilities	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Nondurable Goods	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Leisure and Hospitality	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Wholesale Trade	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Private Service-Providing	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Other	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Manufacturing	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Mining and Logging	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Transportation and Warehousing	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Goods-Producing	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Retail Trade	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Trade, Transportation, and Utilities	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Durable Goods	0.000
	Average Weekly Hours of Production and Nonsupervisory Employees: Construction	0.000
Employment: Unemployment	Number of Civilians Unemployed for 27 Weeks and Over	0.025
	Number of Civilians Unemployed for 15 to 26 Weeks	0.015
	Number of Civilians Unemployed for 5 to 14 Weeks	0.010
	Civilian Unemployment Rate	0.008
	Number of Civilians Unemployed for Less Than 5 Weeks	0.002
Exchange Rates	Percentage Change in US / U.K. Foreign Exchange Rate	0.007
	Percentage Change in US / Australia Foreign Exchange Rate	0.006
	Percentage Change in Canada / US Foreign Exchange Rate	0.003

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Exchange Rates	Percentage Change in China / US Foreign Exchange Rate	0.001
	Percentage Change in Japan / US Foreign Exchange Rate	0.001
	Percentage Change in Switzerland / US Foreign Exchange Rate	0.000
Housing	Percentage Change in Housing Starts in Midwest Census Region	0.134
	Percentage Change in Housing Starts in Northeast Census Region	0.083
	Percentage Change in Privately Owned Housing Starts: 1-Unit Structures	0.081
	Log Housing Starts: 2-4 Units	0.066
	Percentage Change in Privately Owned Housing Starts: 5-Unit Structures or More	0.013
	Percentage Change in Housing Starts in West Census Region	0.012
	Percentage Change in Housing Starts in South Census Region	0.005
Interest Rates	3-Month Treasury Bill: Secondary Market Rate	0.078
	6-Month Treasury Bill: Secondary Market Rate	0.075
	1-Year Treasury Constant Maturity Rate	0.072
	5-Year Treasury Constant Maturity Rate	0.036
	10-Year Treasury Constant Maturity Rate	0.018
Money Supply	Percentage Change in Total Reserves of Depository Institutions	0.069
	Percentage Change in Commercial and Industrial Loans, All Commercial Banks	0.04
	Percentage Change in Monetary Base; Total	0.037
	Percentage Change in M2 Money Stock	0.001
	Percentage Change in M1 Money Stock	0.001
	Percentage Change in Consumer Loans, All Commercial Banks	0.000
Petroleum: Downstream	Percentage Change in US Total Gasoline Retail Sales by Refiners (Thousand Gallons per Day)	0.163
	Percentage Change in US Product Supplied of Finished Petroleum Products (Thousand Barrels)	0.034
	Percentage Change in US Product Supplied of Hydrocarbon Gas Liquids (Thousand Barrels)	0.019
	Percentage Change in US Product Supplied of Crude Oil and Petroleum Products (Thousand Barrels)	0.018
	US Percent Utilization of Refinery Operable Capacity	0.000

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Petroleum: Imports & Exports	Percentage Change in US Exports of Finished Petroleum Products (Thousand Barrels)	0.020
	Percentage Change in US Imports from Non-OPEC Countries of Total Petroleum Products (Thousand Barrels)	0.015
	Percentage Change in US Imports from OPEC Countries of Crude Oil (Thousand Barrels)	0.011
	Percentage Change in US Exports of Crude Oil (Thousand Barrels)	0.003
	Percentage Change in US Imports from Persian Gulf Countries of Crude Oil (Thousand Barrels)	0.003
	Percentage Change in US Imports from OPEC Countries of Total Petroleum Products (Thousand Barrels)	0.001
	Percentage Change in US Imports from Non-OPEC Countries of Crude Oil (Thousand Barrels)	0.001
Petroleum: Prices	Percentage Change in New York Harbor No. 2 Heating Oil Future Contract 3 (Dollars per Gallon)	0.044
	Percentage Change in Cushing, OK Crude Oil Future Contract 4 (Dollars per Barrel)	0.021
	Percentage Change in Cushing, OK Crude Oil Future Contract 3 (Dollars per Barrel)	0.020
	Percentage Change in Cushing, OK Crude Oil Future Contract 2 (Dollars per Barrel)	0.019
	Percentage Change in Cushing, OK Crude Oil Future Contract 1 (Dollars per Barrel)	0.018
	Percentage Change in New York Harbor No. 2 Heating Oil Future Contract 1 (Dollars per Gallon)	0.017
	Percentage Change in US Gulf Coast Kerosene-Type Jet Fuel Spot Price FOB (Dollars per Gallon)	0.010
	Percentage Change in New York Harbor Conventional Gasoline Regular Spot Price FOB (Dollars per Gallon)	0.009
	Percentage Change in New York Harbor No. 2 Heating Oil Spot Price FOB (Dollars per Gallon)	0.008
	Percentage Change in US Gulf Coast Conventional Gasoline Regular Spot Price FOB (Dollars per Gallon)	0.006
	Percentage Change in Mont Belvieu, TX Propane Spot Price FOB (Dollars per Gallon)	0.003
Petroleum: Storage	Percentage Change in US Ending Stocks of Hydrocarbon Gas Liquids (Thousand Barrels)	0.176
	Percentage Change in US Ending Stocks of Gasoline Blending Components (Thousand Barrels)	0.166
	Percentage Change in US Ending Stocks of Petroleum Coke (Thousand Barrels)	0.054

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Petroleum: Storage	Percentage Change in US Ending Stocks of Asphalt and Road Oil (Thousand Barrels)	0.052
	Percentage Change in US Ending Stocks of Unfinished Oils (Thousand Barrels)	0.029
	Percentage Change in US Ending Stocks of Conventional Motor Gasoline (Thousand Barrels)	0.025
	Percentage Change in US Crude Oil Stocks at Refineries (Thousand Barrels)	0.021
	Percentage Change in US Ending Stocks of Finished Motor Gasoline (Thousand Barrels)	0.019
	Percentage Change in US Ending Stocks of Total Petroleum Products (Thousand Barrels)	0.013
	Percentage Change in US Ending Stocks of Crude Oil in SPR (Thousand Barrels)	0.013
	Percentage Change in US Ending Stocks of Miscellaneous Petroleum Products (Thousand Barrels)	0.012
	Percentage Change in US Crude Oil Stocks at Tank Farms and Pipelines (Thousand Barrels)	0.010
	Percentage Change in US Ending Stocks of Distillate Fuel Oil (Thousand Barrels)	0.007
	Percentage Change in US Ending Stocks of Lubricants (Thousand Barrels)	0.006
	Percentage Change in US Ending Stocks of Kerosene-Type Jet Fuel (Thousand Barrels)	0.004
	Percentage Change in US Ending Stocks of Aviation Gasoline (Thousand Barrels)	0.003
	Percentage Change in US Ending Stocks of Waxes (Thousand Barrels)	0.001
	Percentage Change in US Crude Oil Stocks in Transit (on Ships) from Alaska (Thousand Barrels)	0.000
	Percentage Change in US Ending Stocks of Kerosene (Thousand Barrels)	0.000
	Percentage Change in US Ending Stocks of Residual Fuel Oil (Thousand Barrels)	0.000
Petroleum: Upstream	Percentage Change in US Crude Oil Rotary Rigs in Operation	0.025
	Percentage Change in US Field Production of Crude Oil (Thousand Barrels)	0.000
Producer Price Indices	Percentage Change in Producer Price Index by Industry: Drilling Oil and Gas Wells	0.025
	Percentage Change in Producer Price Index by Industry: Total Manufacturing Industries	0.023
	Percentage Change in Producer Price Index for All Commodities	0.019
	Percentage Change in Producer Price Index by Industry: Support Activities for Oil and Gas Operations	0.008
	Percentage Change in Producer Price Index by Industry: Oil and Gas Field Machinery and Equipment Manufacturing	0.004

Table 3.3 Continued

Category	Transformed Series Name	R-Squared
Producer Price Indices	Percentage Change in Producer Price Index by Industry: Total Mining Industries	0.003
Production Indices	Percentage Change in Industrial Production: Nonindustrial supplies	0.629
	Percentage Change in Industrial Production: Miscellaneous durable goods	0.619
	Percentage Change in Industrial Production: General business supplies	0.571
	Percentage Change in Industrial Production: Construction supplies	0.531
	Percentage Change in Industrial Production: Manufacturing (SIC)	0.435
	Percentage Change in Industrial Production: Business supplies	0.402
	Percentage Change in Industrial Production: Final Products and Nonindustrial Supplies	0.379
	Percentage Change in Industrial Production: Durable manufacturing	0.354
	Percentage Change in Industrial Production: Other equipment parts	0.353
	Percentage Change in Industrial Production: Computers, communications equipment, and	0.353
	Percentage Change in Industrial Production: Equipment parts	0.353
	Percentage Change in Industrial Production: Computer and other board assemblies and parts	0.326
	Percentage Change in Industrial Production: Finished processing	0.304
	Percentage Change in Industrial Production: Industrial and other equipment	0.303
	Percentage Change in Industrial Production: Miscellaneous nondurable goods	0.299
	Percentage Change in Industrial Production: Final products	0.247
	Percentage Change in Industrial Production: Other equipment	0.239
	Percentage Change in Industrial Production: Industrial equipment	0.236
	Percentage Change in Industrial Production: Other manufacturing	0.215
	Percentage Change in Industrial Production: Miscellaneous nondurable materials	0.180
	Percentage Change in Industrial Production: Durable consumer goods	0.152
	Percentage Change in Industrial Production: Consumer goods	0.146
	Percentage Change in Industrial Production: Other nondurable materials	0.131
	Percentage Change in Industrial Production: Defense and space equipment	0.126
	Percentage Change in Industrial Production Index	0.118
	Percentage Change in Industrial Production: Mining	0.101

Table 3.3 Continued

Category	Transformed Series Names	R-Squared
Production Indices	Percentage Change in Industrial Production: Consumer parts	0.088
	Percentage Change in Industrial Production: Residential utilities	0.063
	Percentage Change in Industrial Production: Consumer energy products	0.058
	Percentage Change in Industrial Production: Electric and gas utilities	0.020

Table 3.4 Median Responses of Selected Variables to the Structural Shocks across All Periods

	Change in Contemporaneous Period			Cumulative Change at One Year		
Variable	Oil Supply	Aggregate Demand	Oil Demand	Oil Supply	Aggregate Demand	Oil Demand
Percentage Change in Oil Prices	3.6*	4.6*	6.2*	23.5*	30.1*	39.6*
Percentage Change in Oil Production	-0.1*	0.4*	0.1*	-1.0*	2.2	0.6
Percentage Change in Oil Storage	-0.3*	-0.7*	-0.6*	-2.0*	-4.5*	-4.2*
Change in Refinery Capacity Utilization (Percentage Points)	-1.0*	0.7*	-0.6*	-6.7*	5.6*	-3.5
Percentage Change in Gasoline Sales	-1.5*	1.2*	-0.7*	-9.9*	8.1*	-4.2
Percentage Change in Industrial Production Index	-0.2*	0.2*	-0.0*	-1.1*	1.2*	-0.3
Percentage Change in Producer Price Index	0.2*	0.4*	0.5*	1.5*	2.7*	3.0*
Percentage Change in Canada-US Exchange Rate	-0.4*	-0.4*	-0.6*	-2.4*	-2.9*	-3.9*

* The entry's 90-percent HPD interval does not contain zero.

Figure 3.1 Plot of oil prices (Cushing WTI Spot Price) over time with recessions shaded

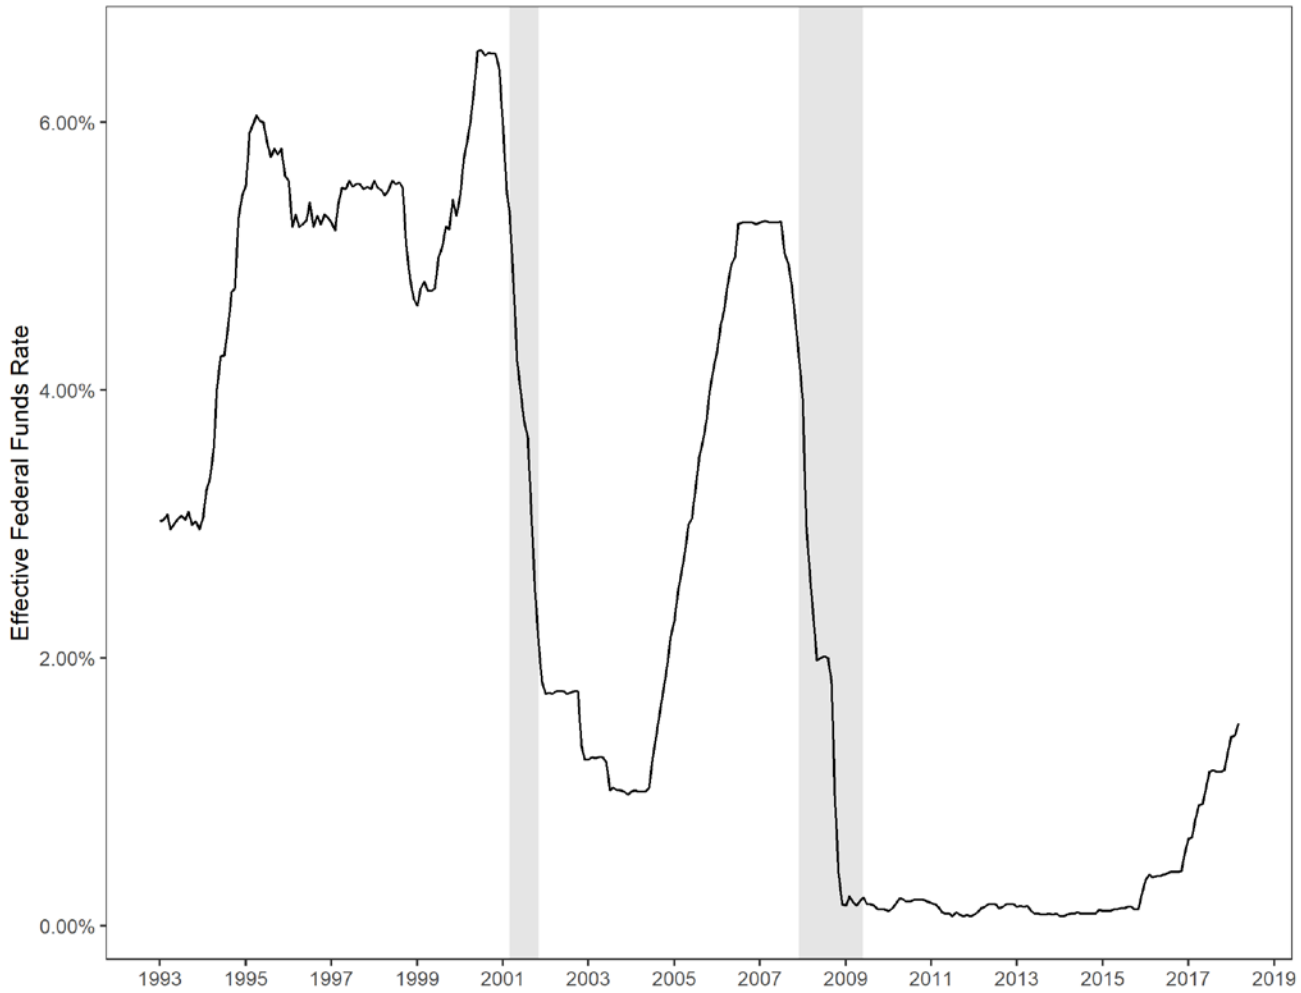

Figure 3.2 Plot of the federal funds rate over time with recessions shaded

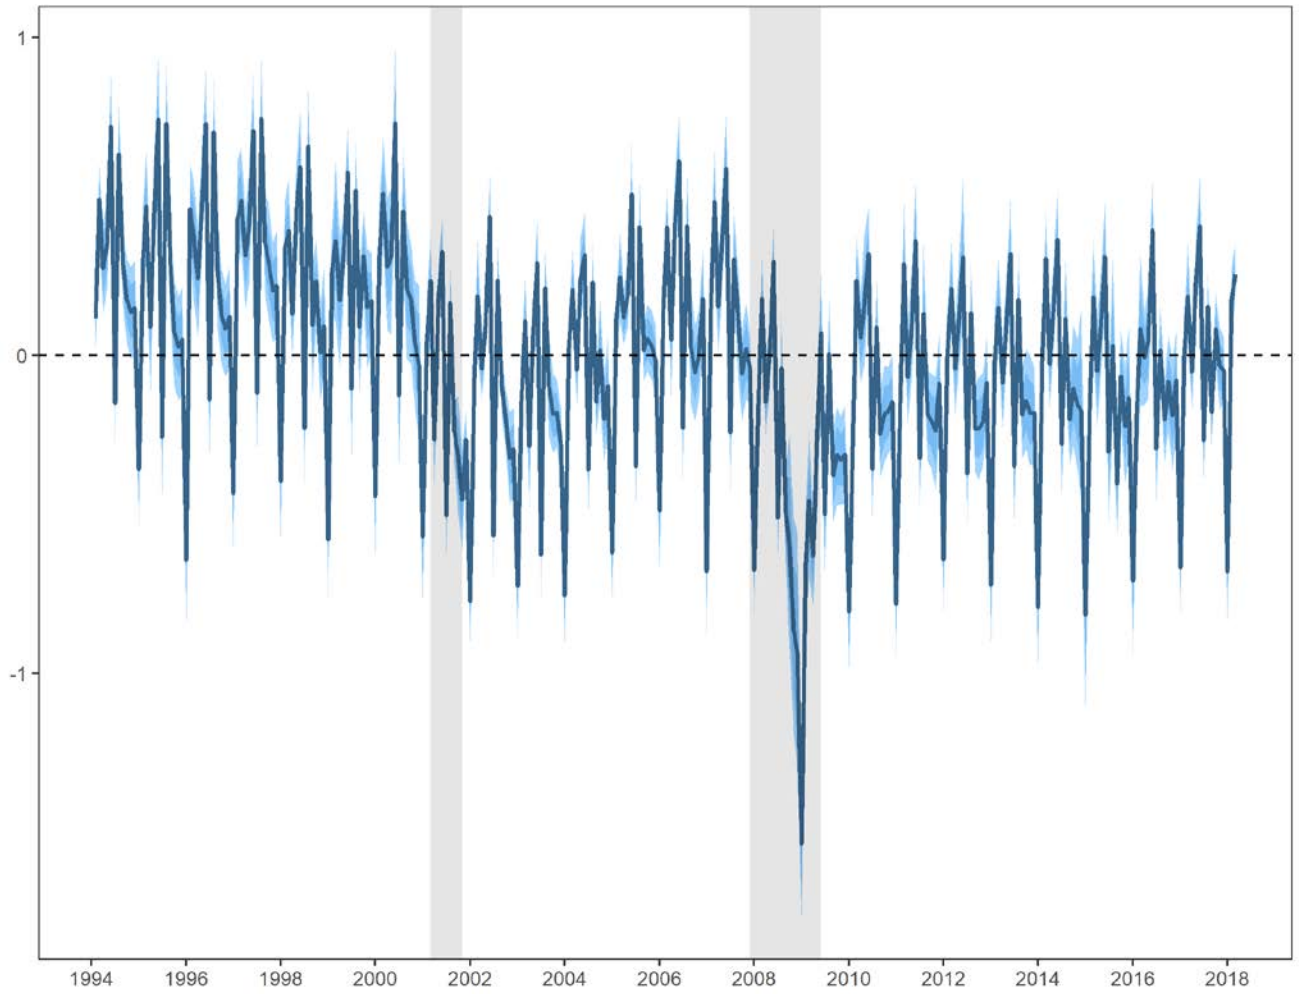

Figure 3.3 Plot of the median of the latent factor with 68- and 90-percent highest posterior densities

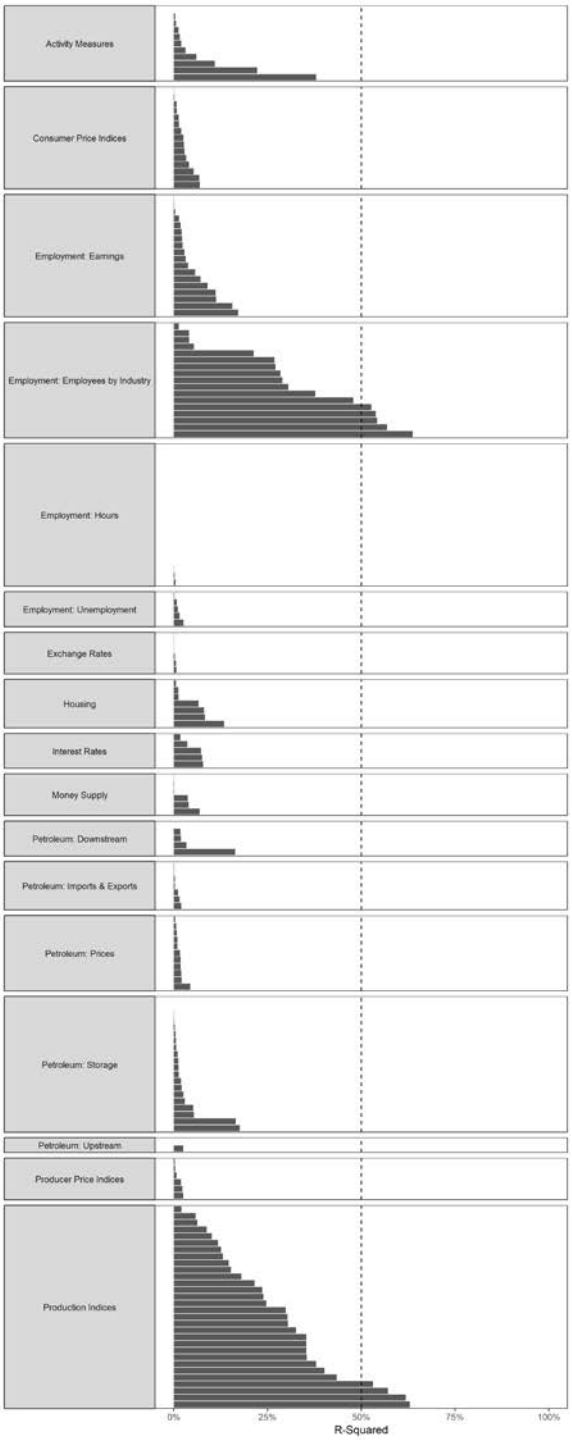

Figure 3.4 Median R-squared values for the informational series (individually) regressed on the latent factor regressed across all MCMC iterations, organized by category

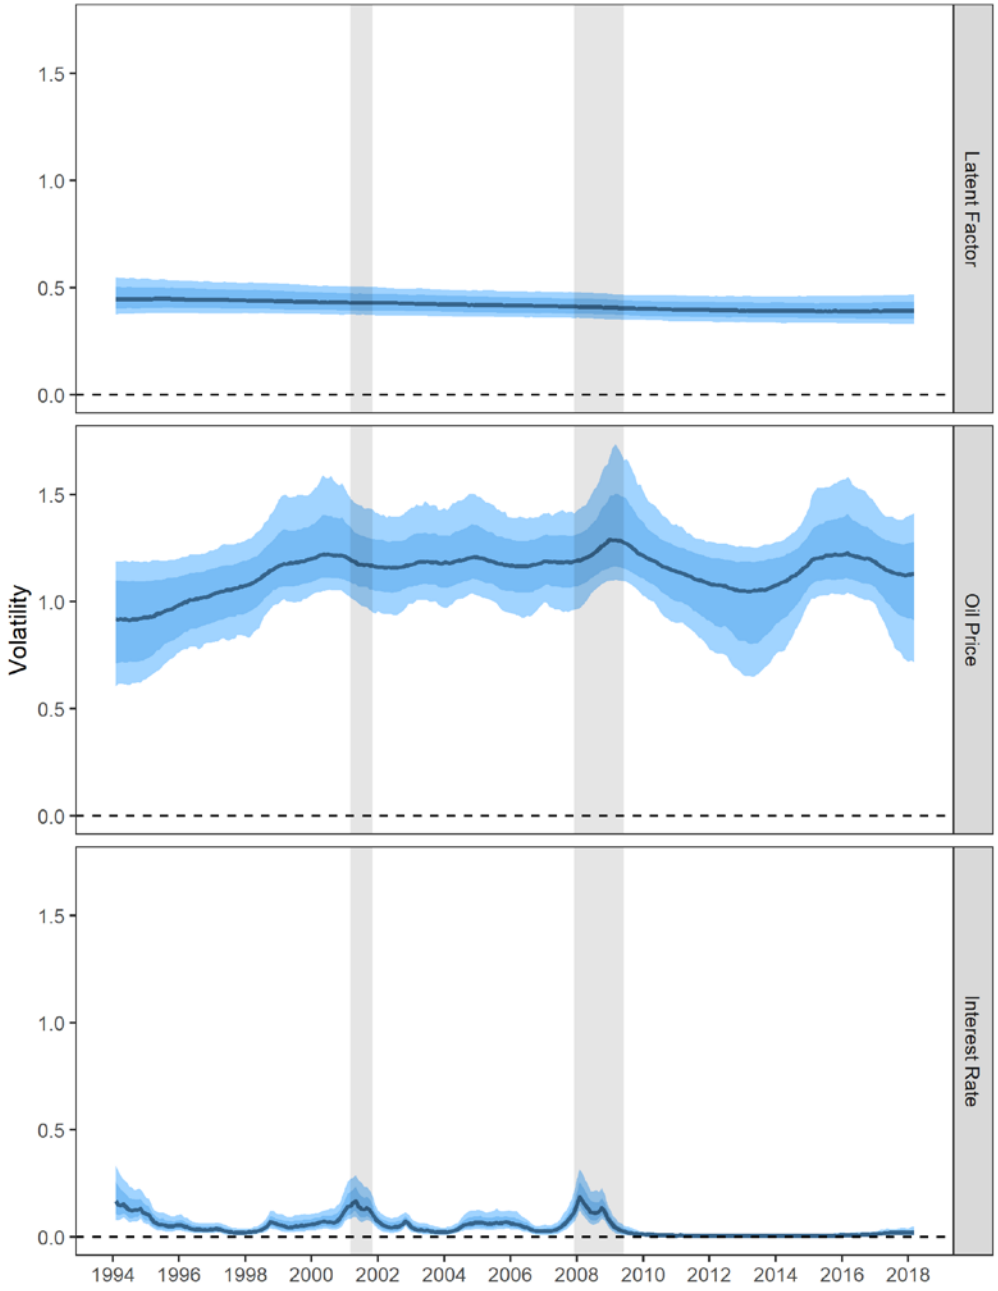

Figure 3.5 Plots of the median of the time-varying residual volatilities for latent factor, oil prices, and federal funds rate with 68- and 90-percent highest posterior densities

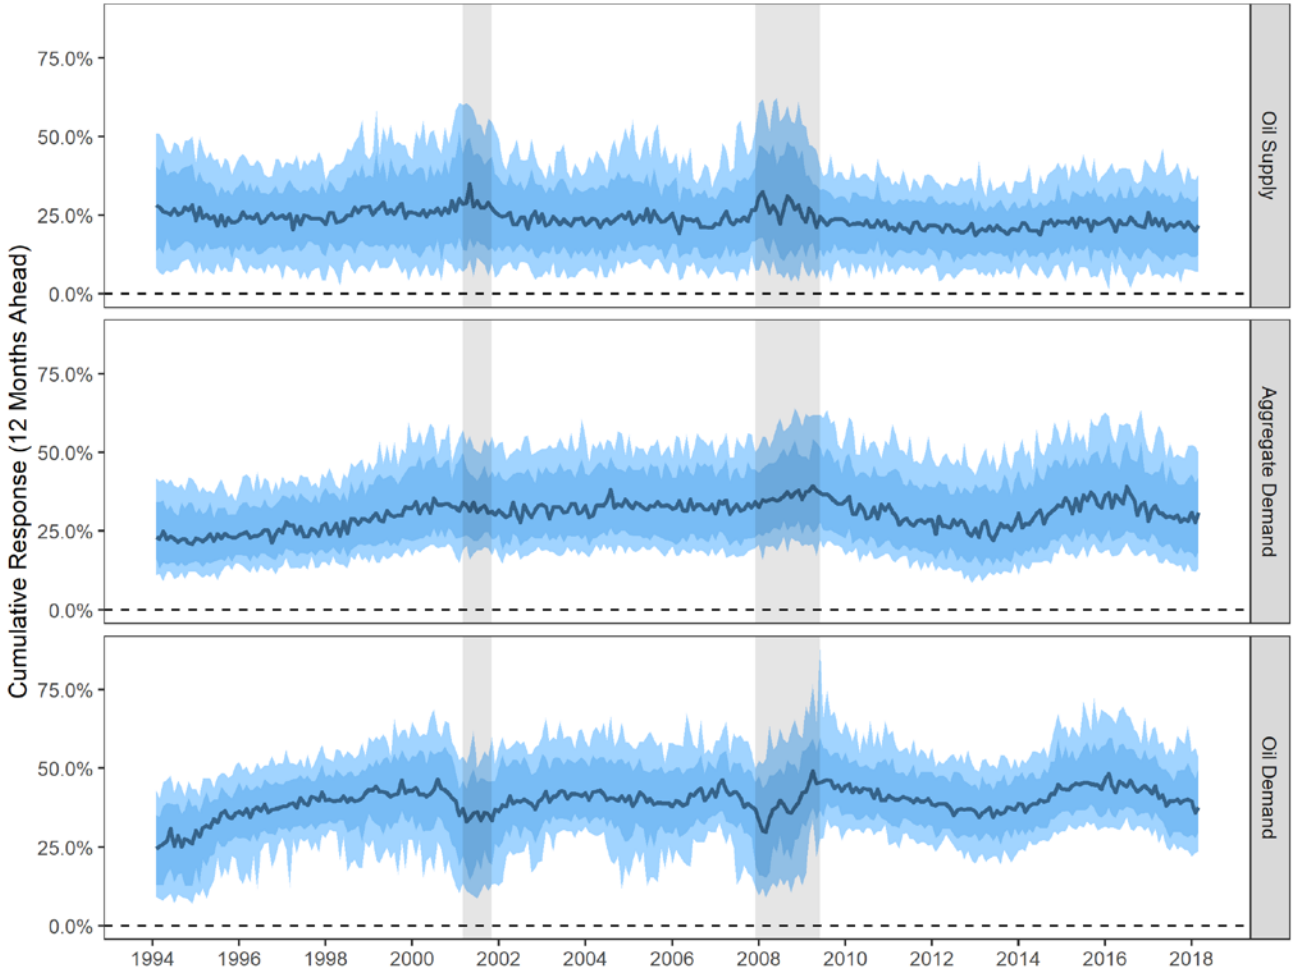

Figure 3.6 Evolution of 12-month cumulative responses of oil prices to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

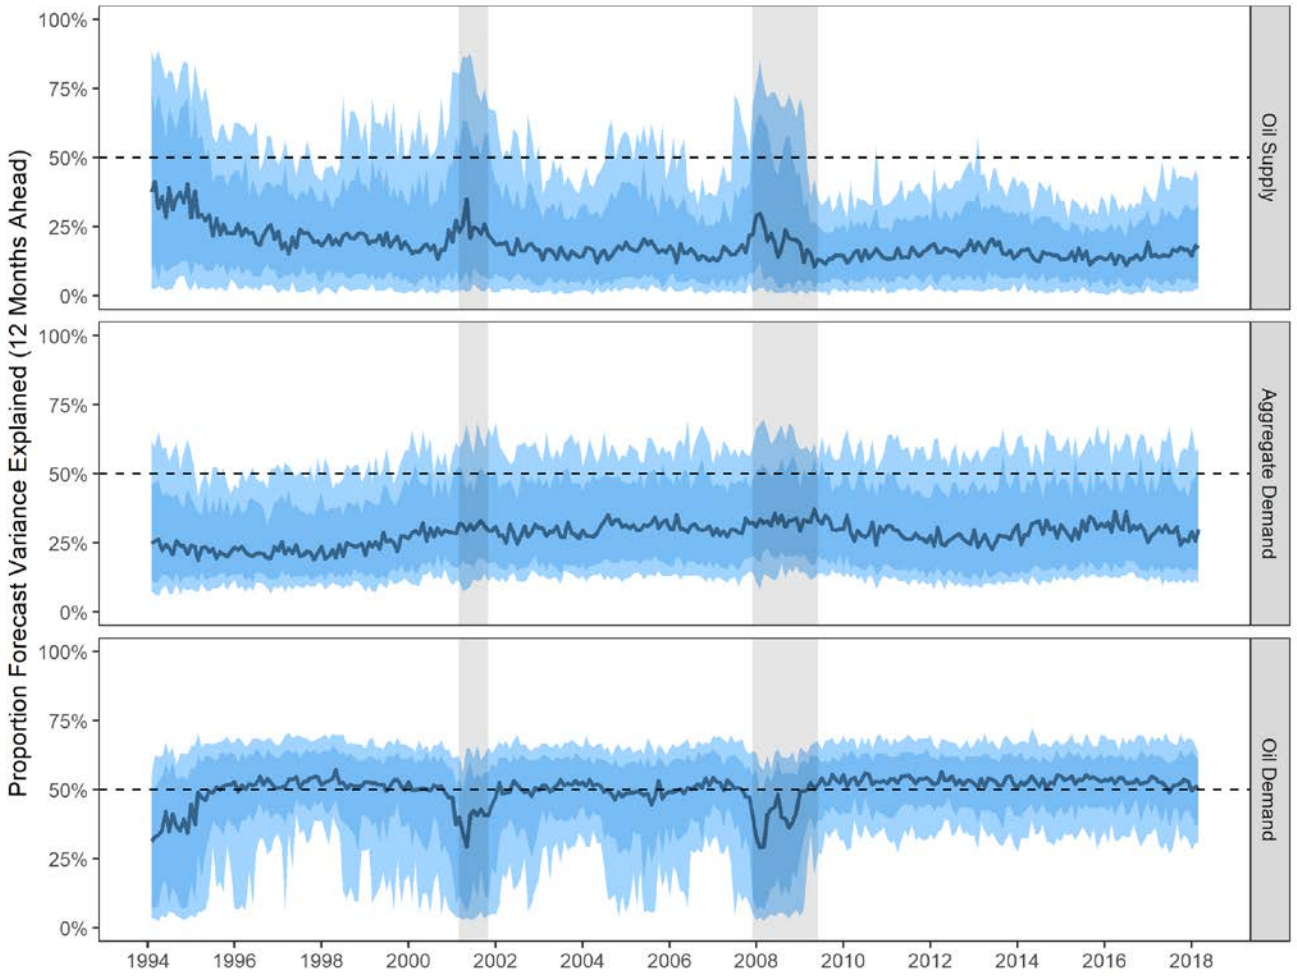

Figure 3.7 Evolution of proportions of the 12-month forecast error variance for oil prices explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

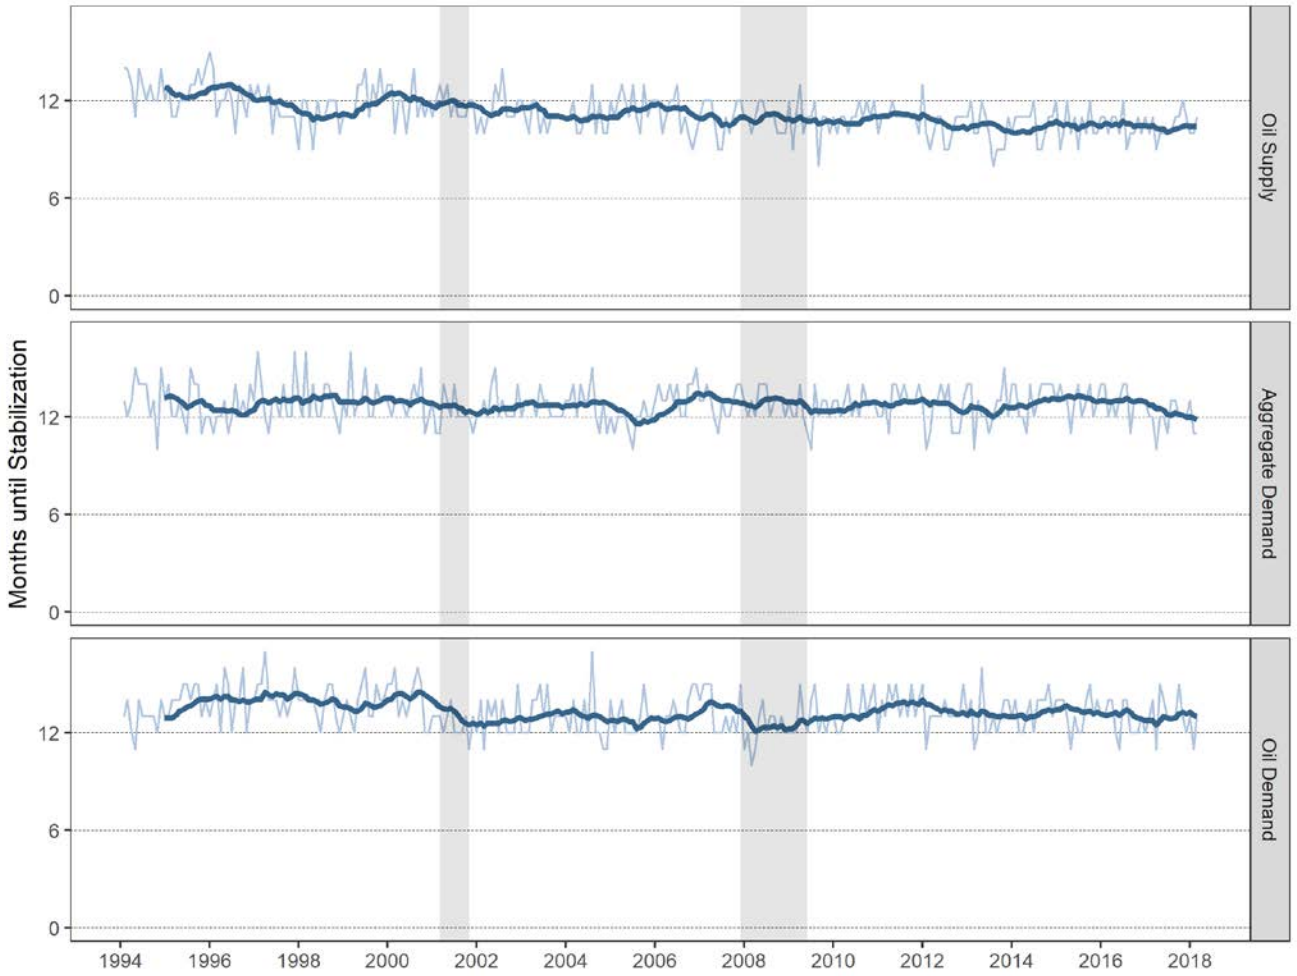

Figure 3.8 Number of months until stabilization of oil prices following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

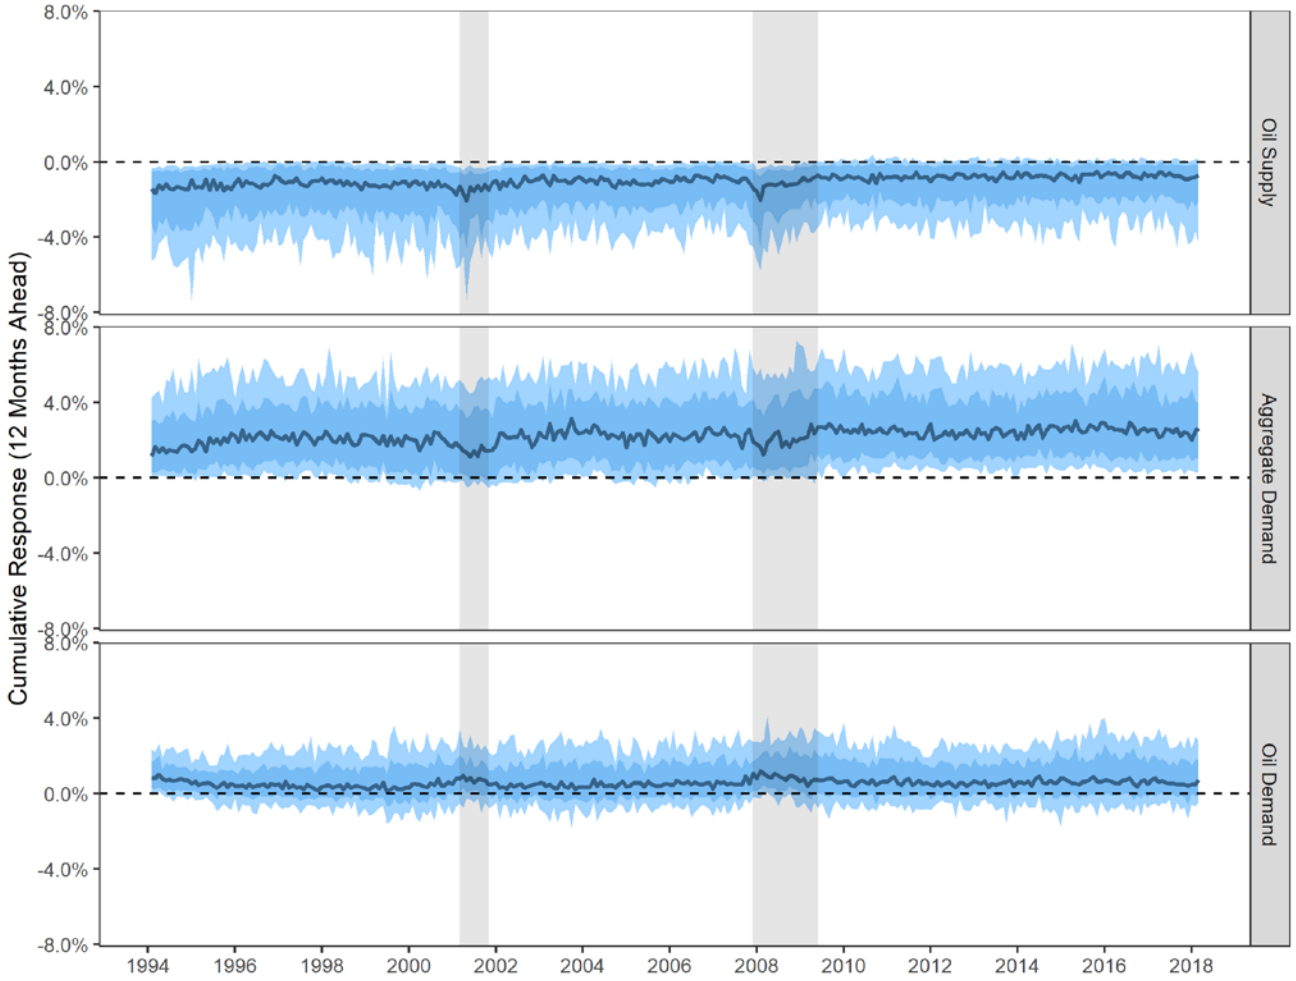

Figure 3.9 Evolution of 12-month cumulative responses of oil production to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

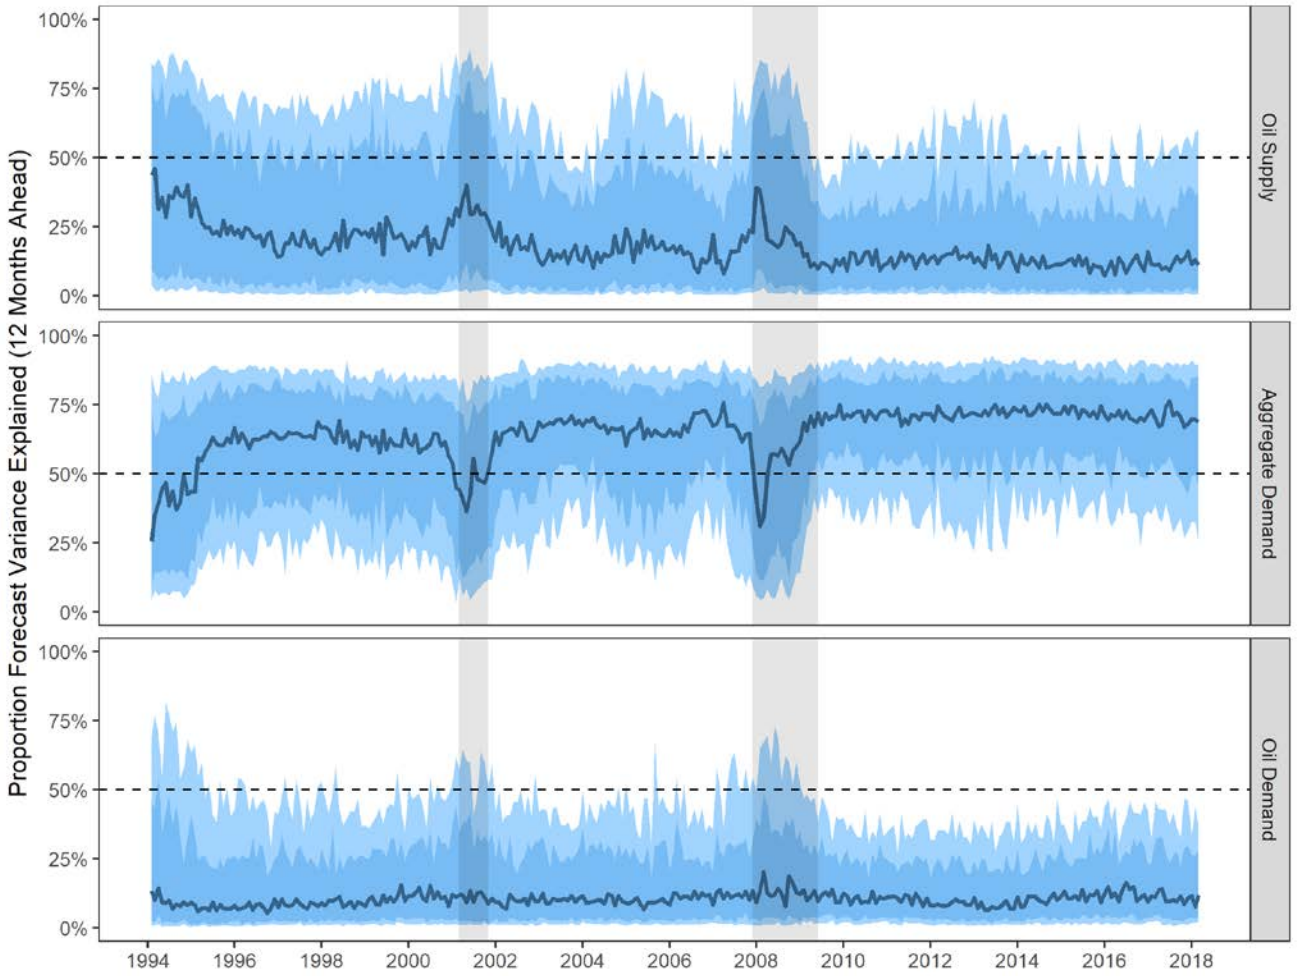

Figure 3.10 Evolution of proportions of the 12-month forecast error variance for oil production explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

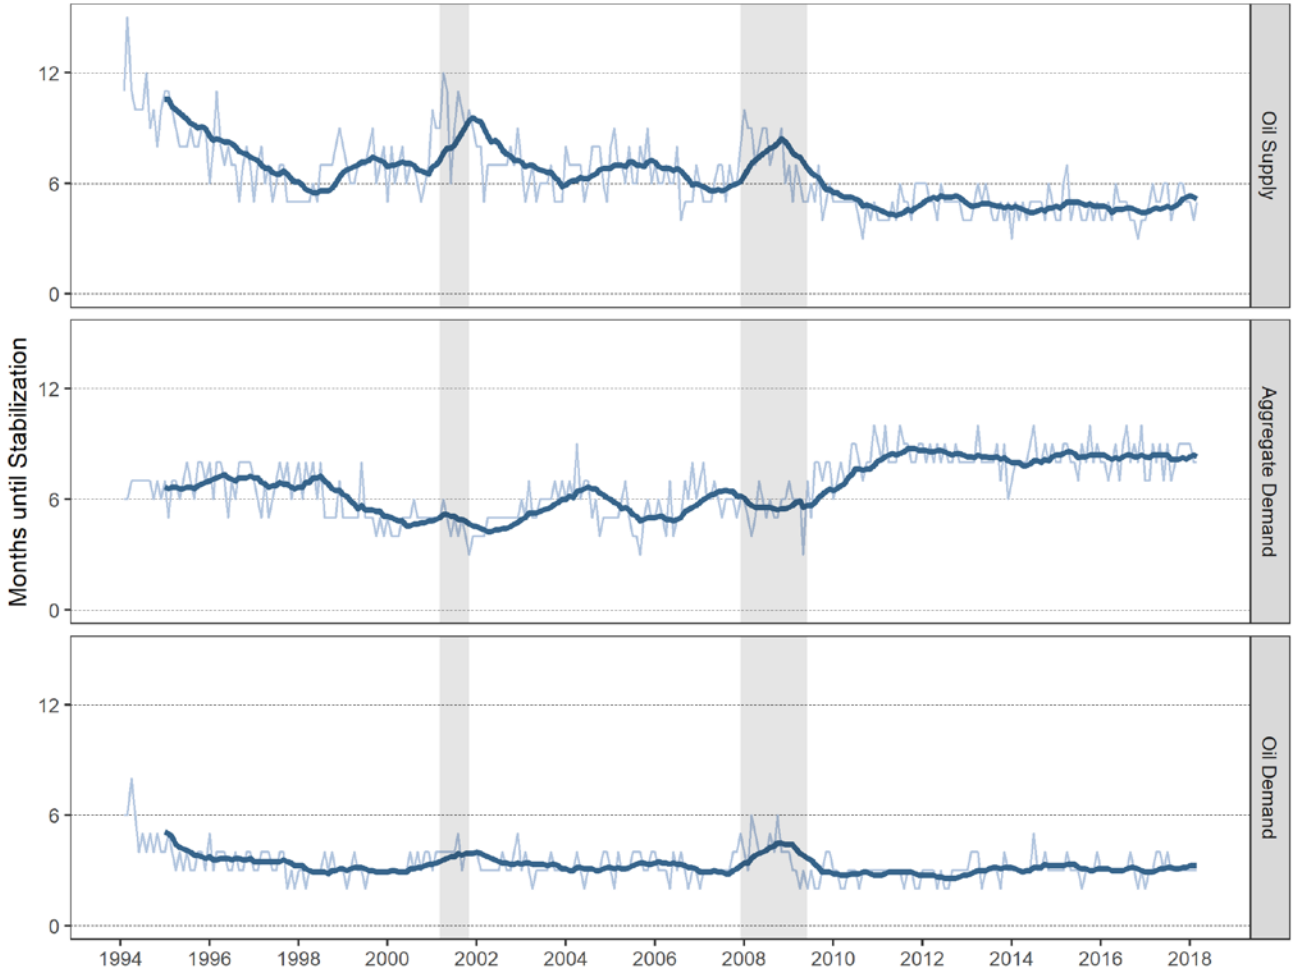

Figure 3.11 Number of months until stabilization of oil production following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

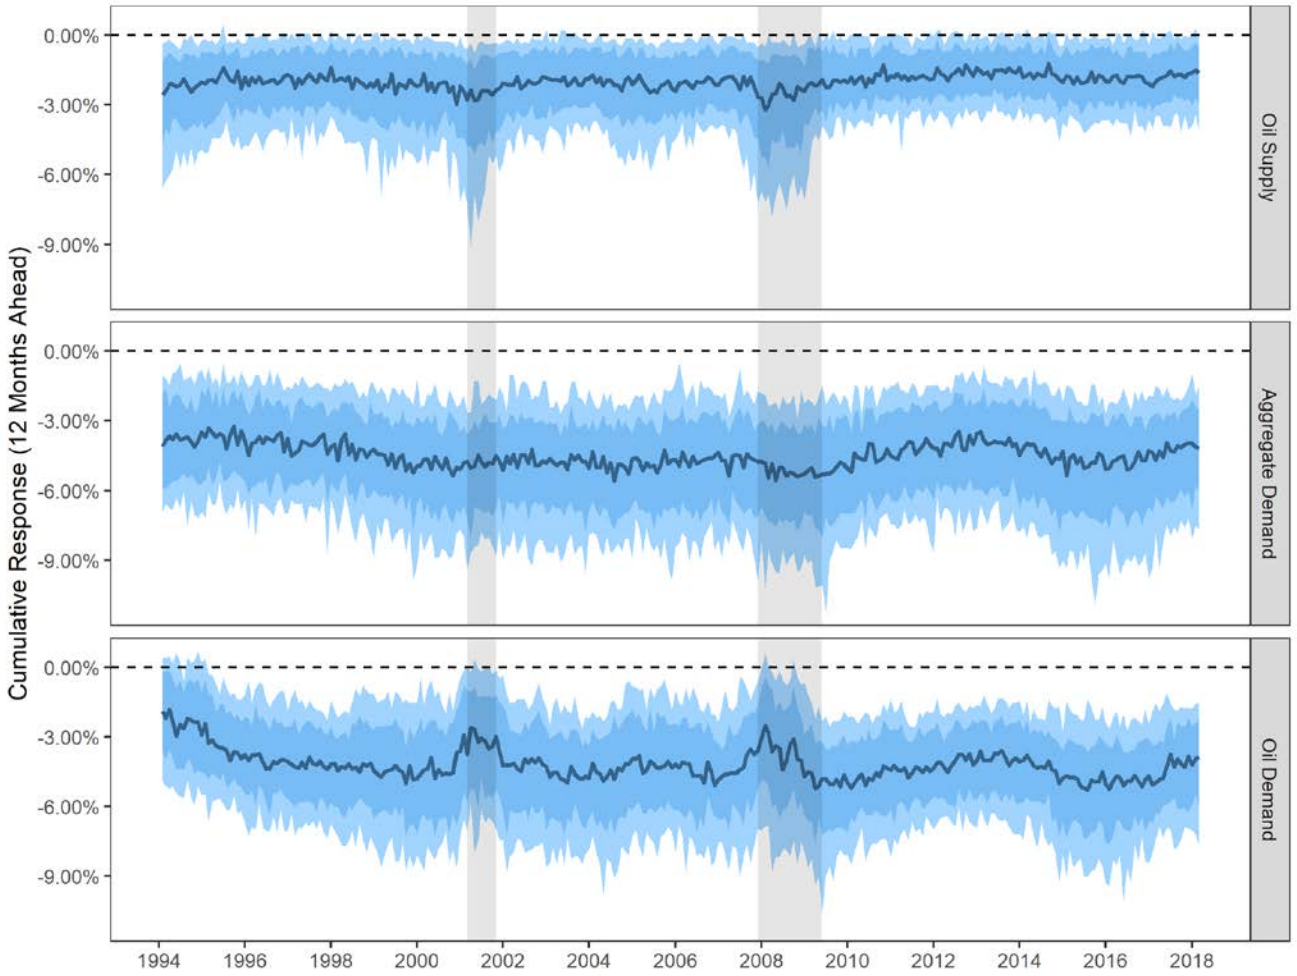

Figure 3.12 Evolution of 12-month cumulative responses of oil storage to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

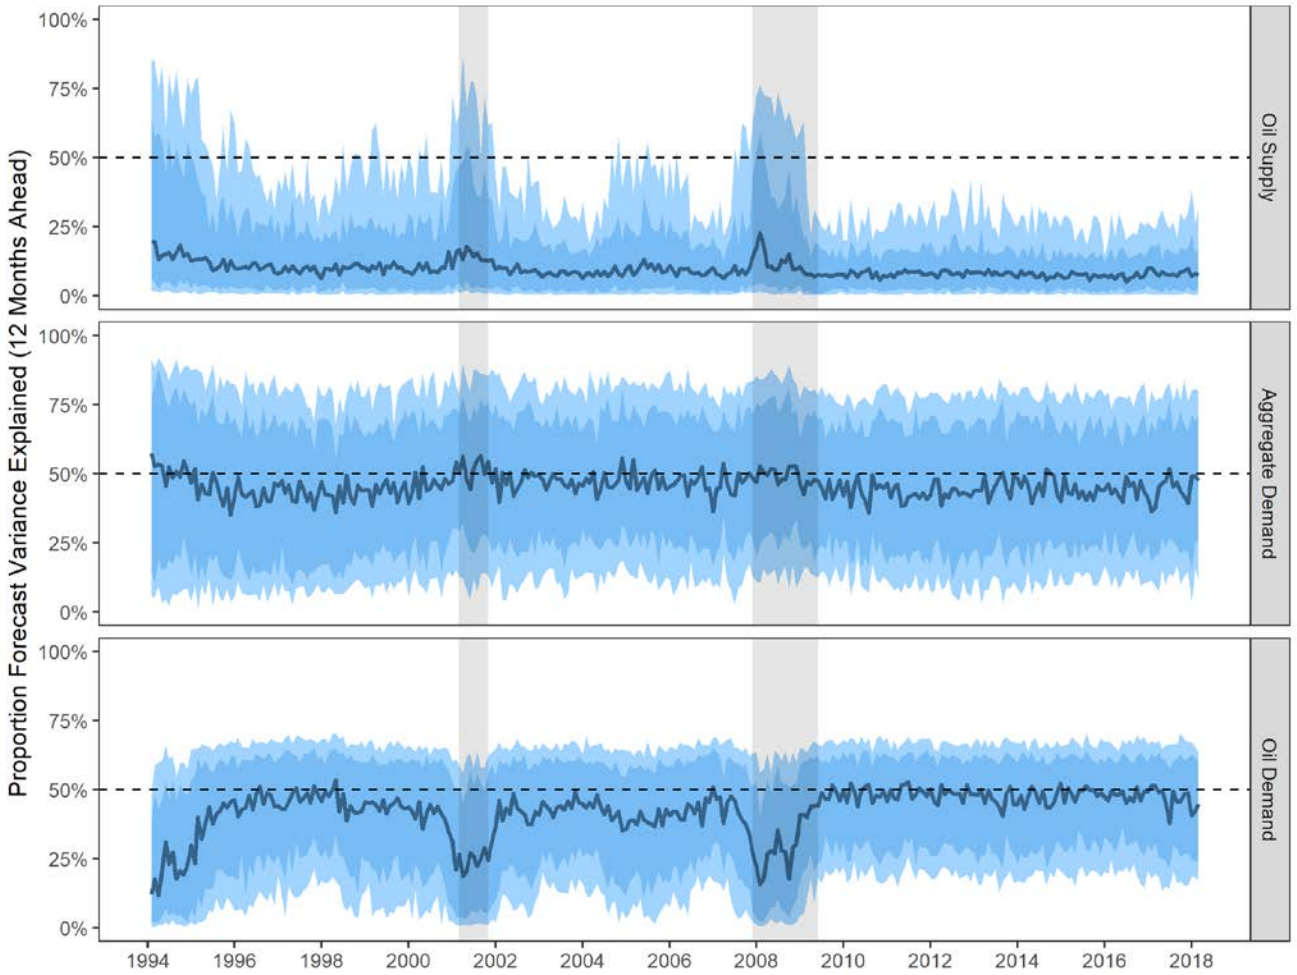

Figure 3.13 Evolution of proportions of the 12-month forecast error variance for oil storage explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

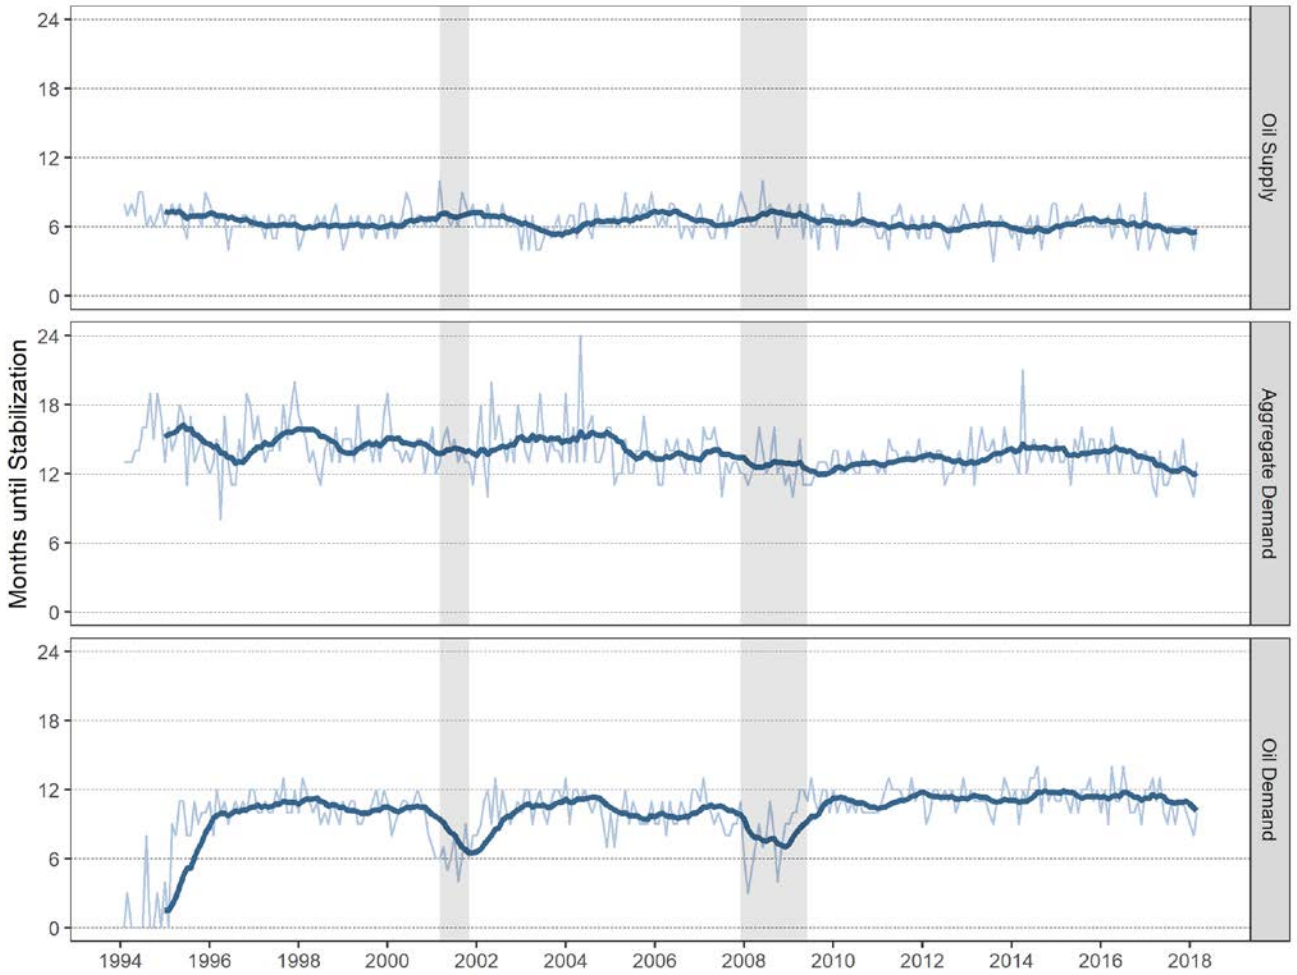

Figure 3.14 Number of months until stabilization of oil storage following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

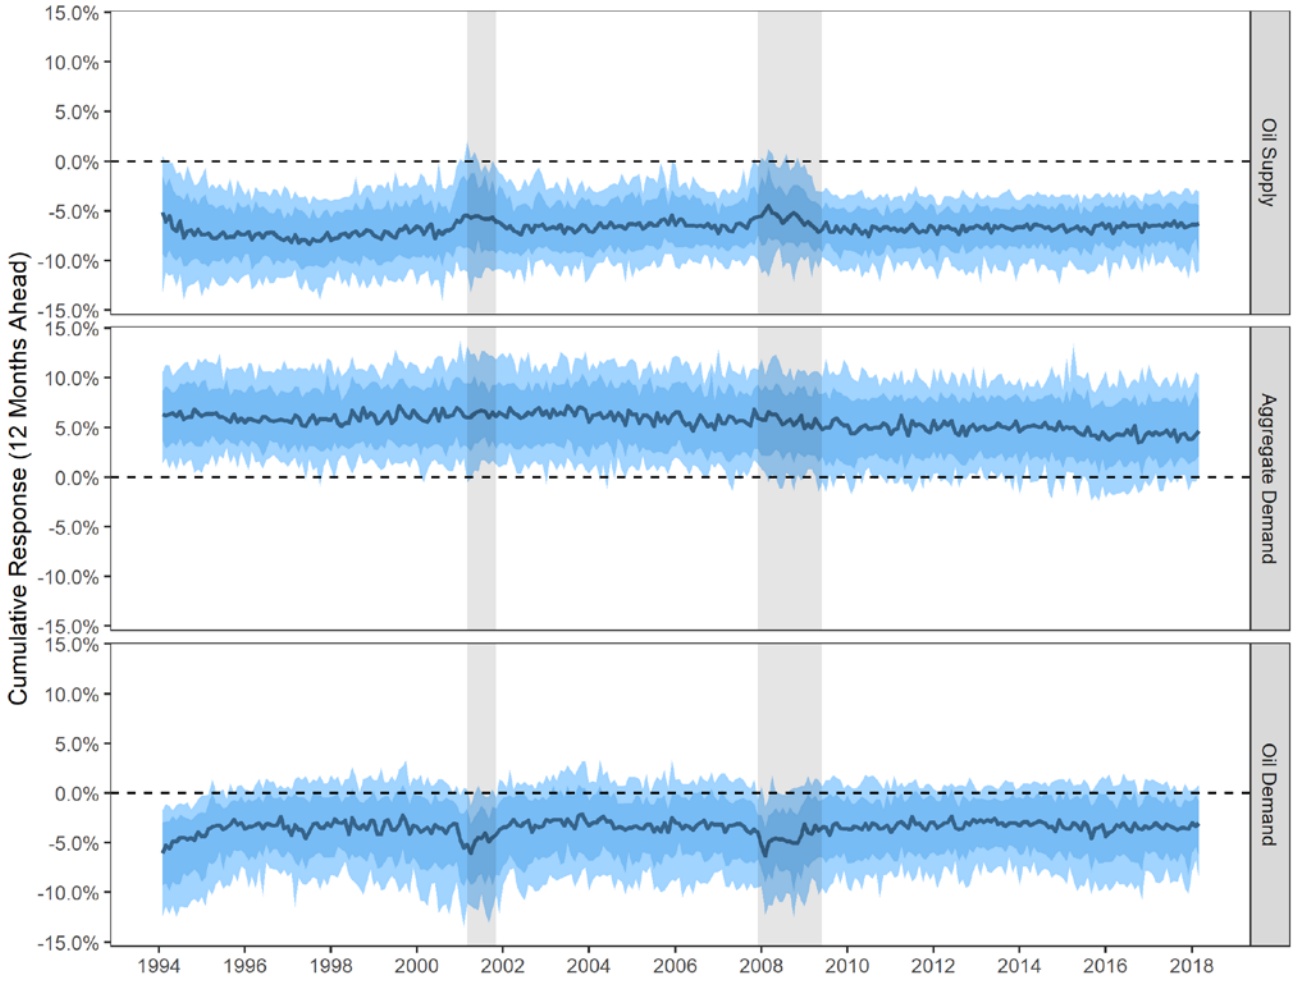

Figure 3.15 Evolution of 12-month cumulative responses of refinery capacity utilization to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

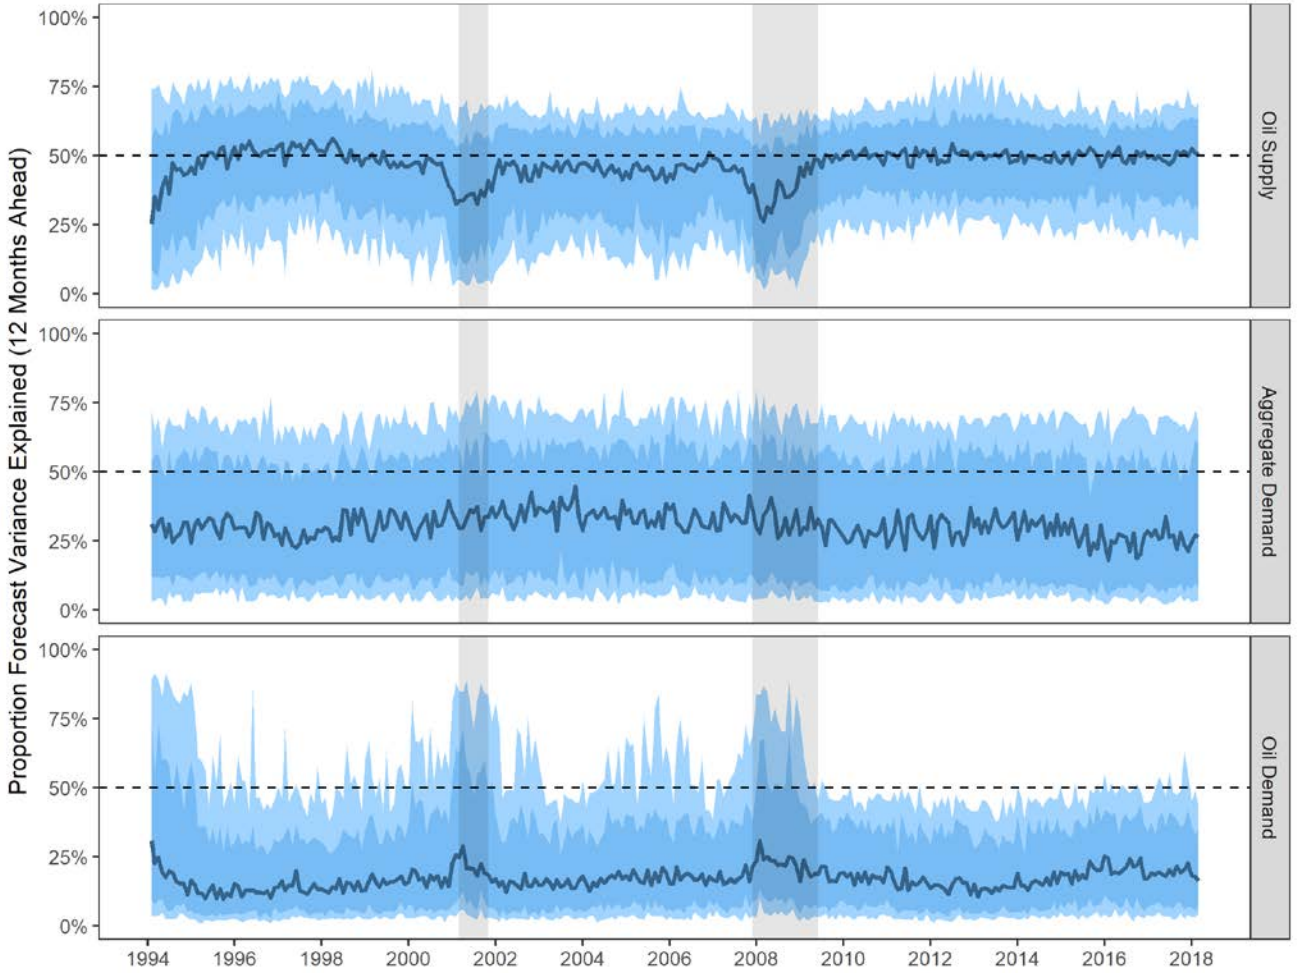

Figure 3.16 Evolution of proportions of the 12-month forecast error variance for refinery capacity utilization explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

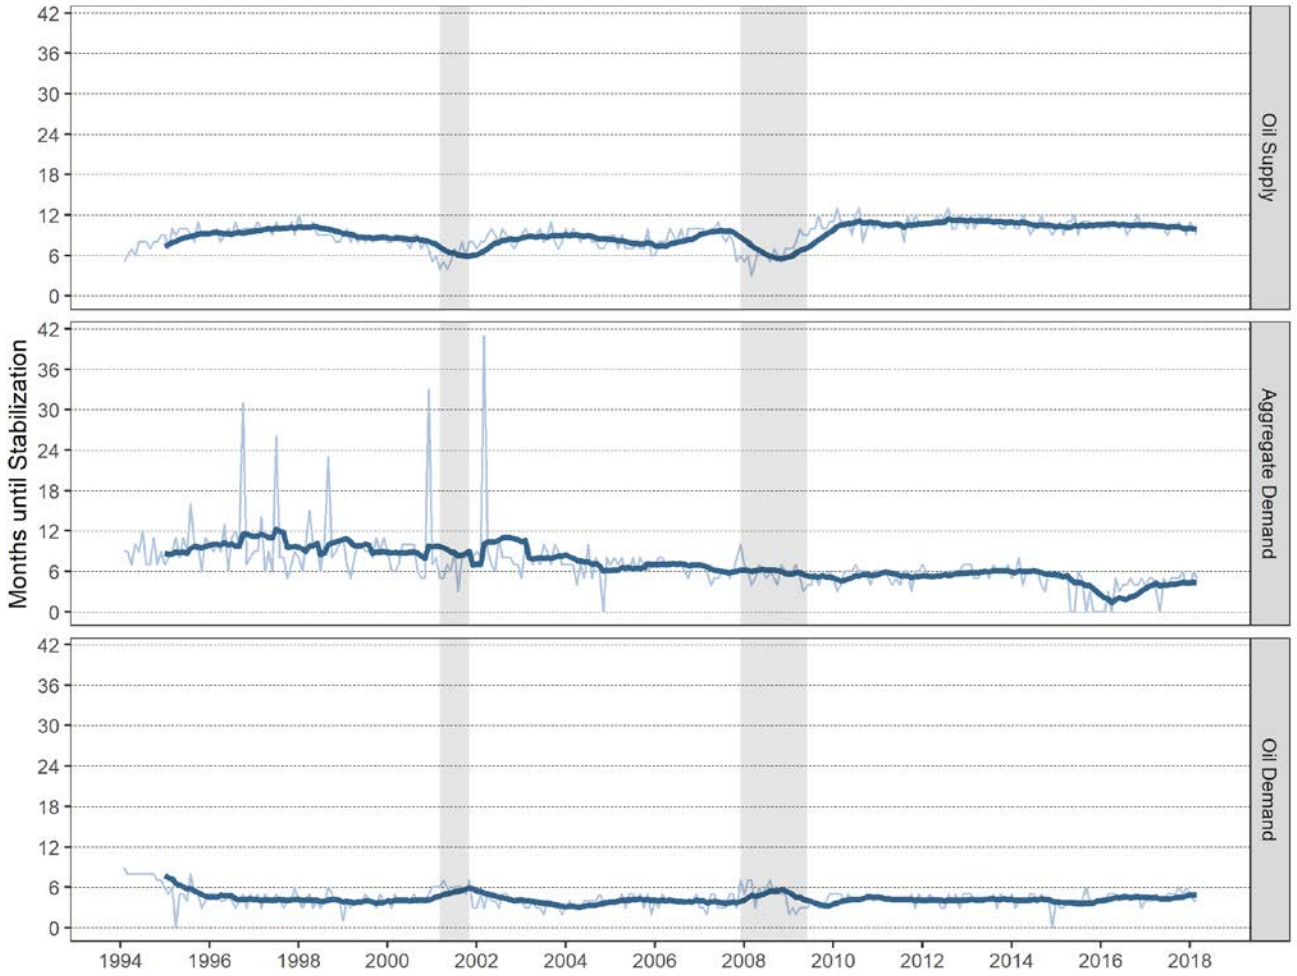

Figure 3.17 Number of months until stabilization of refinery capacity utilization following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

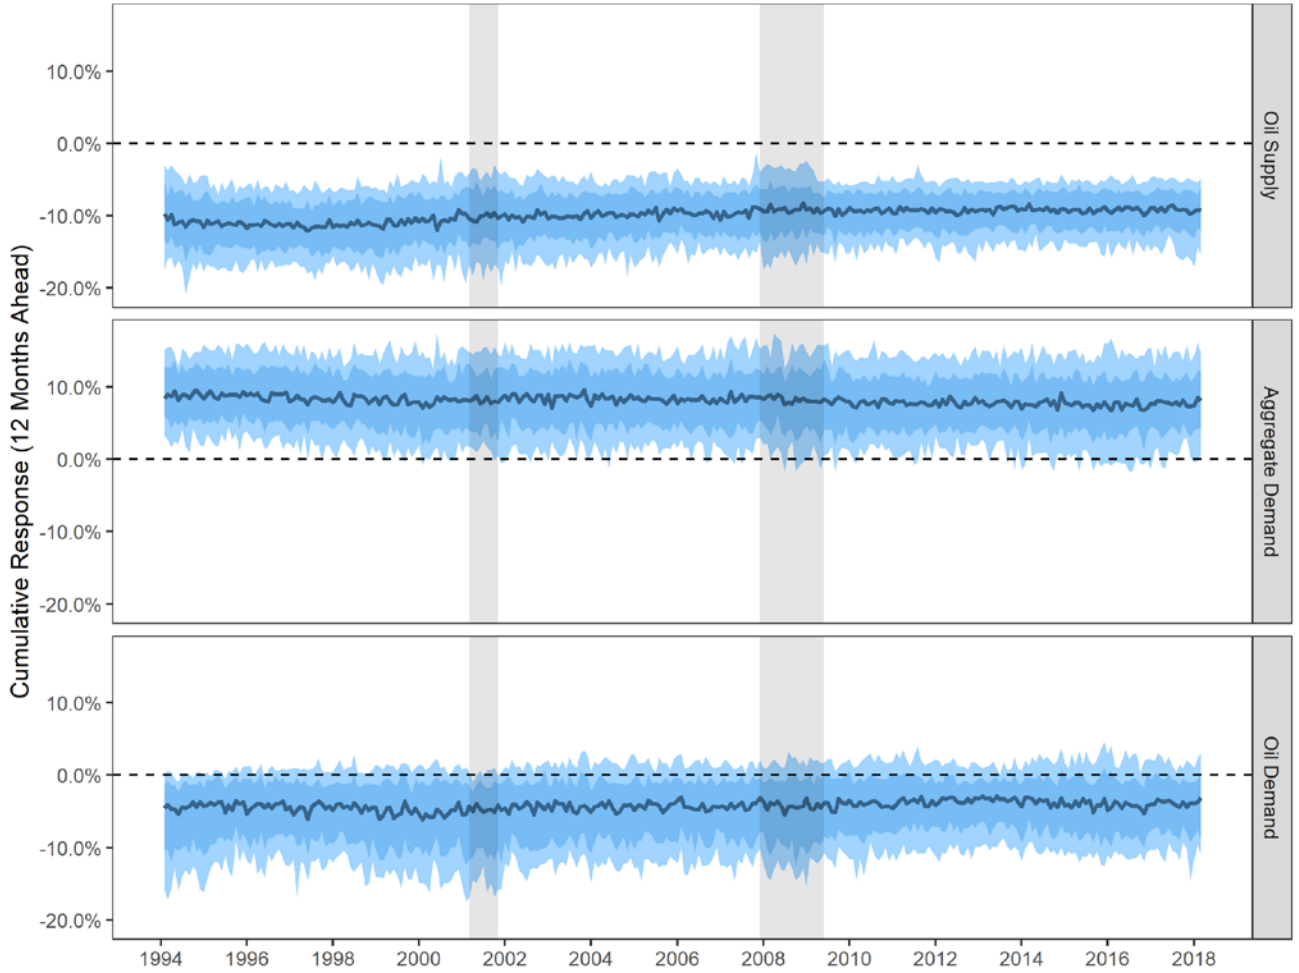

Figure 3.18 Evolution of 12-month cumulative responses of gasoline sales to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

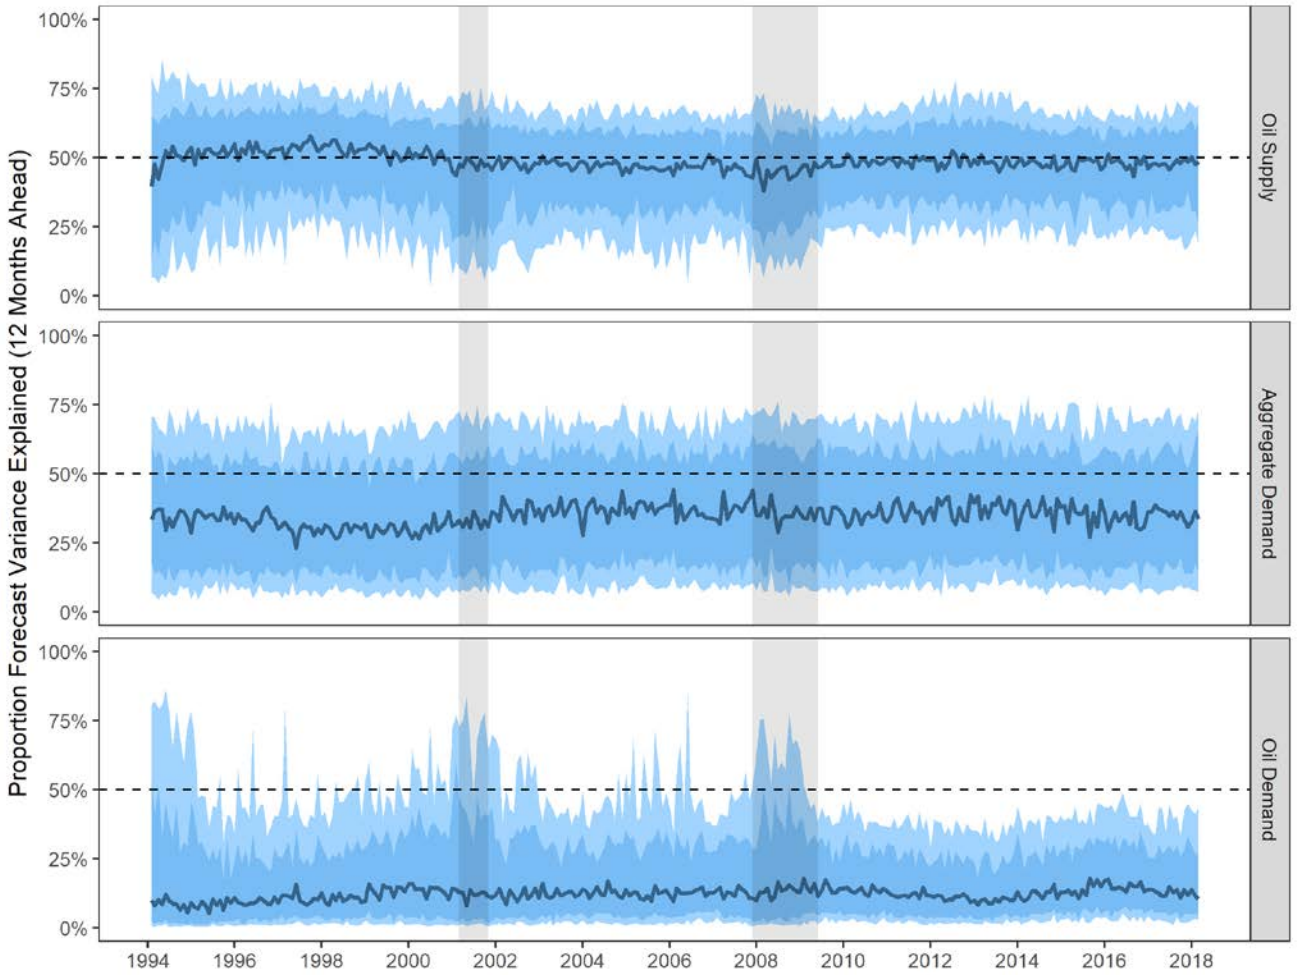

Figure 3.19 Evolution of proportions of the 12-month forecast error variance for gasoline sales explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

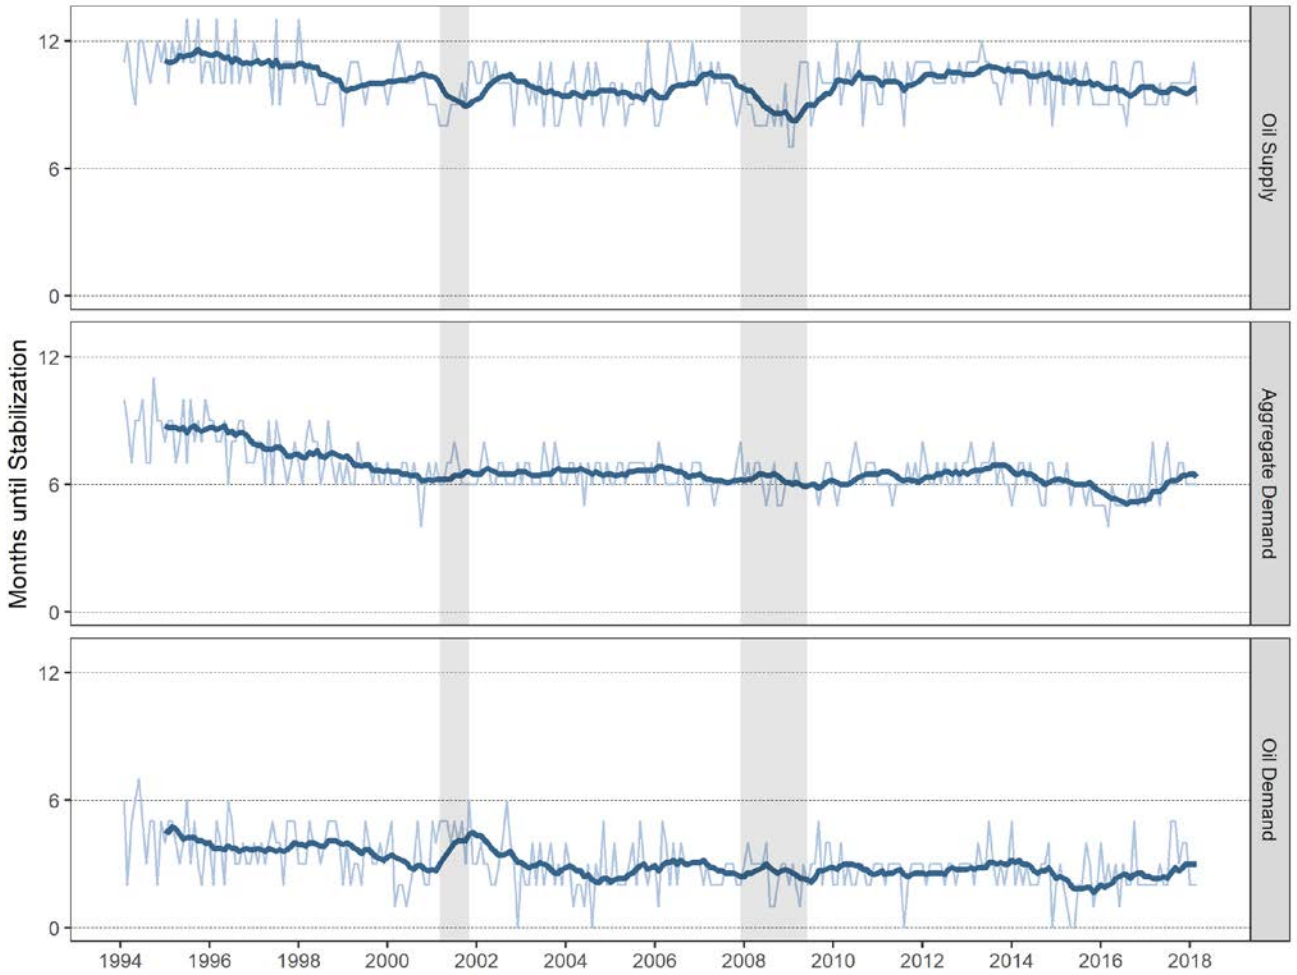

Figure 3.20 Number of months until stabilization of gasoline sales following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

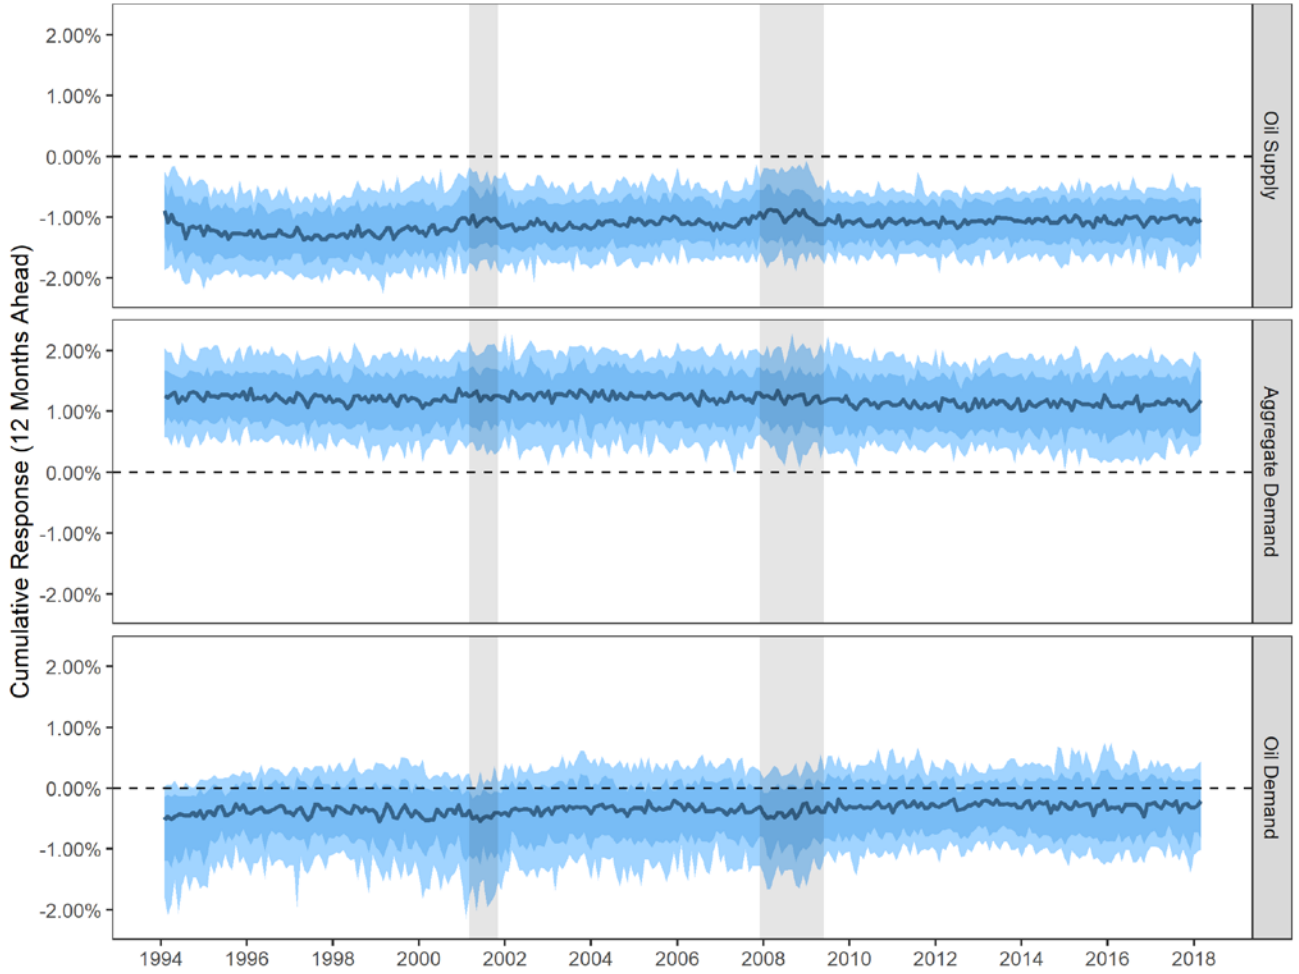

Figure 3.21 Evolution of 12-month cumulative responses of Industrial Production Index to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

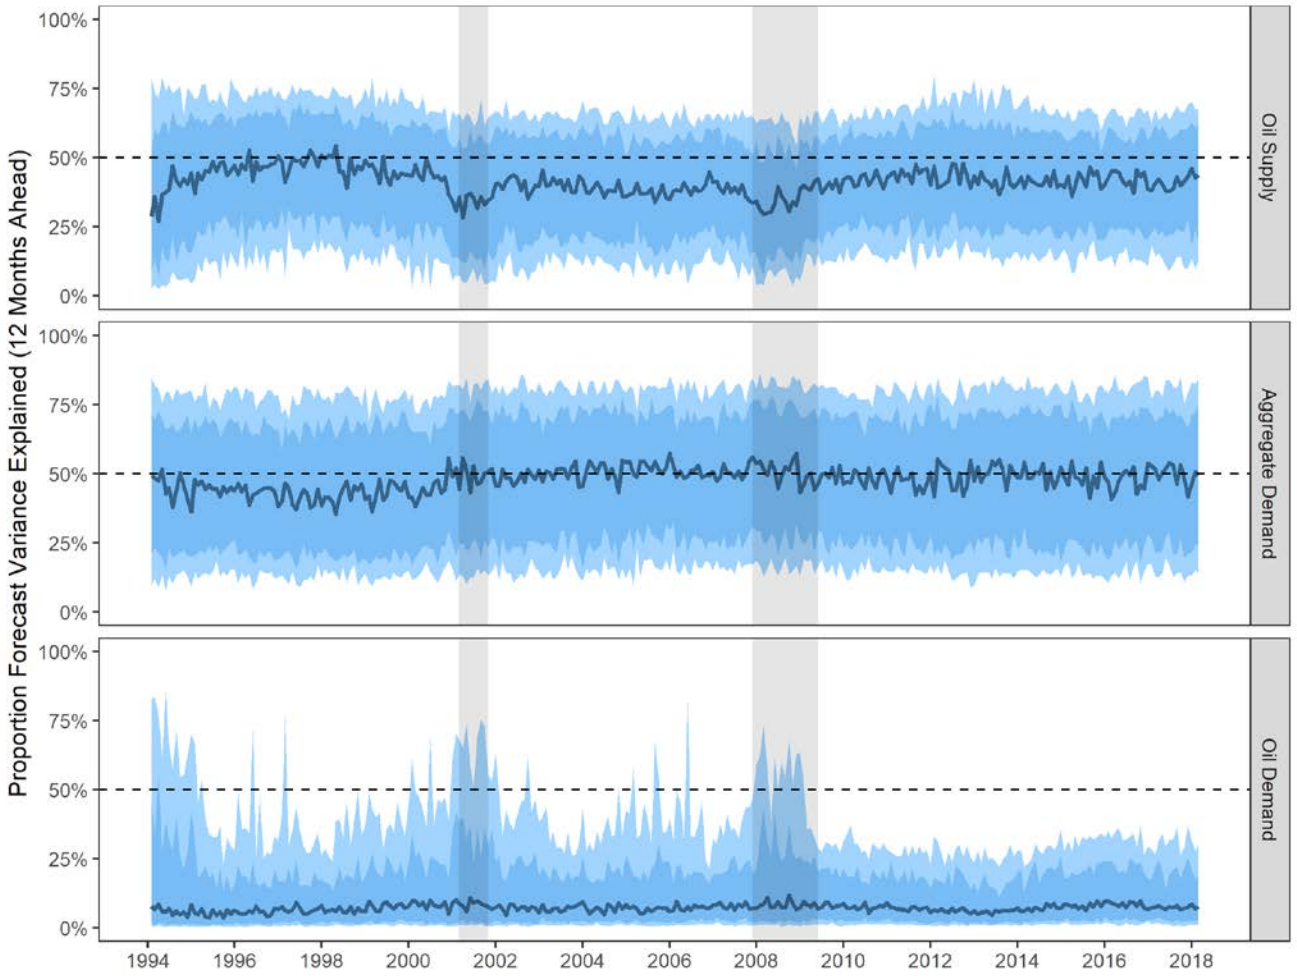

Figure 3.22 Evolution of proportions of the 12-month forecast error variance for Industrial Production Index explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

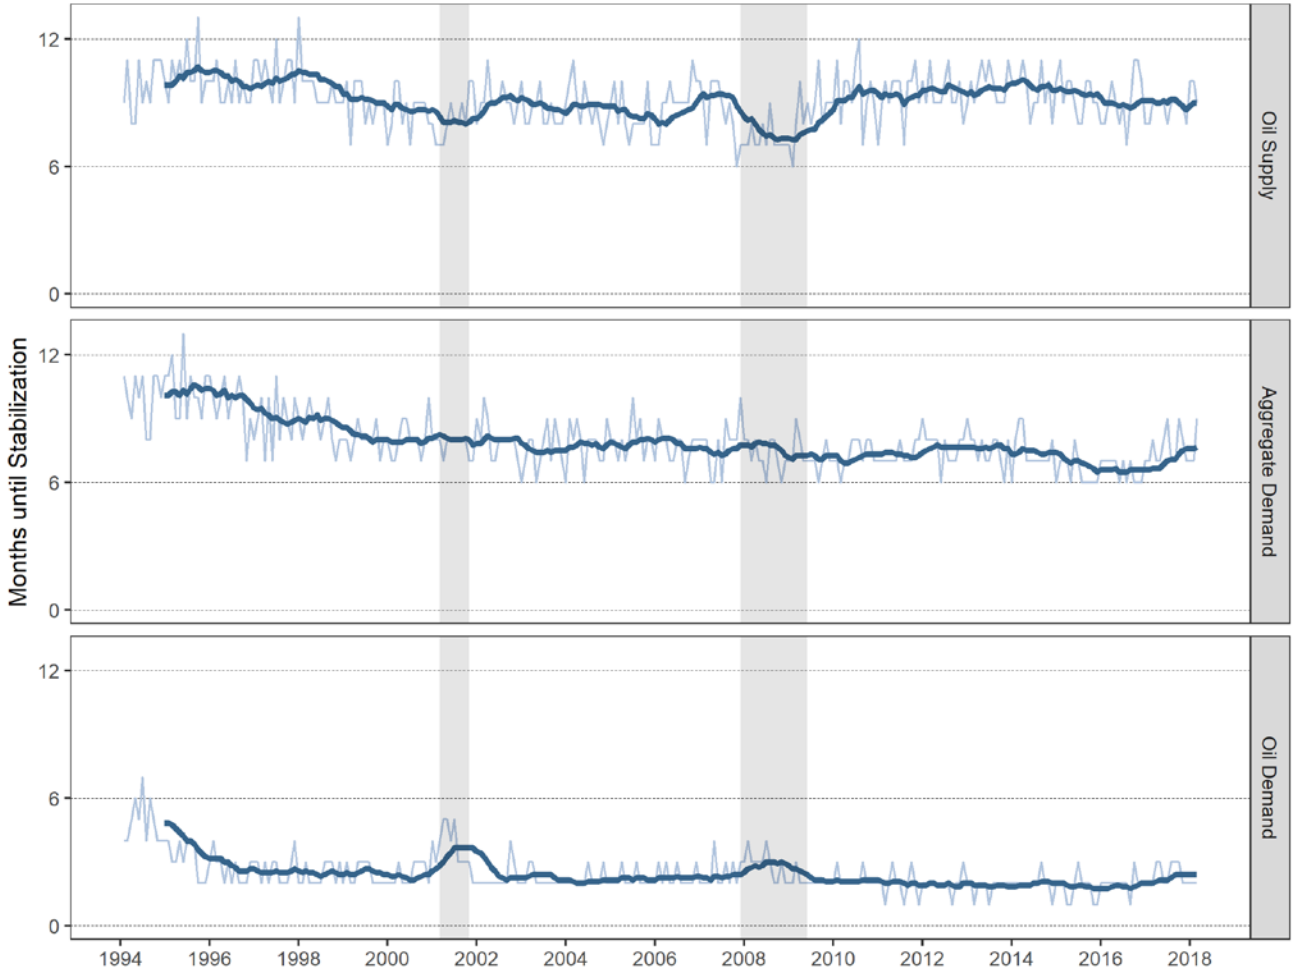

Figure 3.23 Number of months until stabilization of Industrial Production Index following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

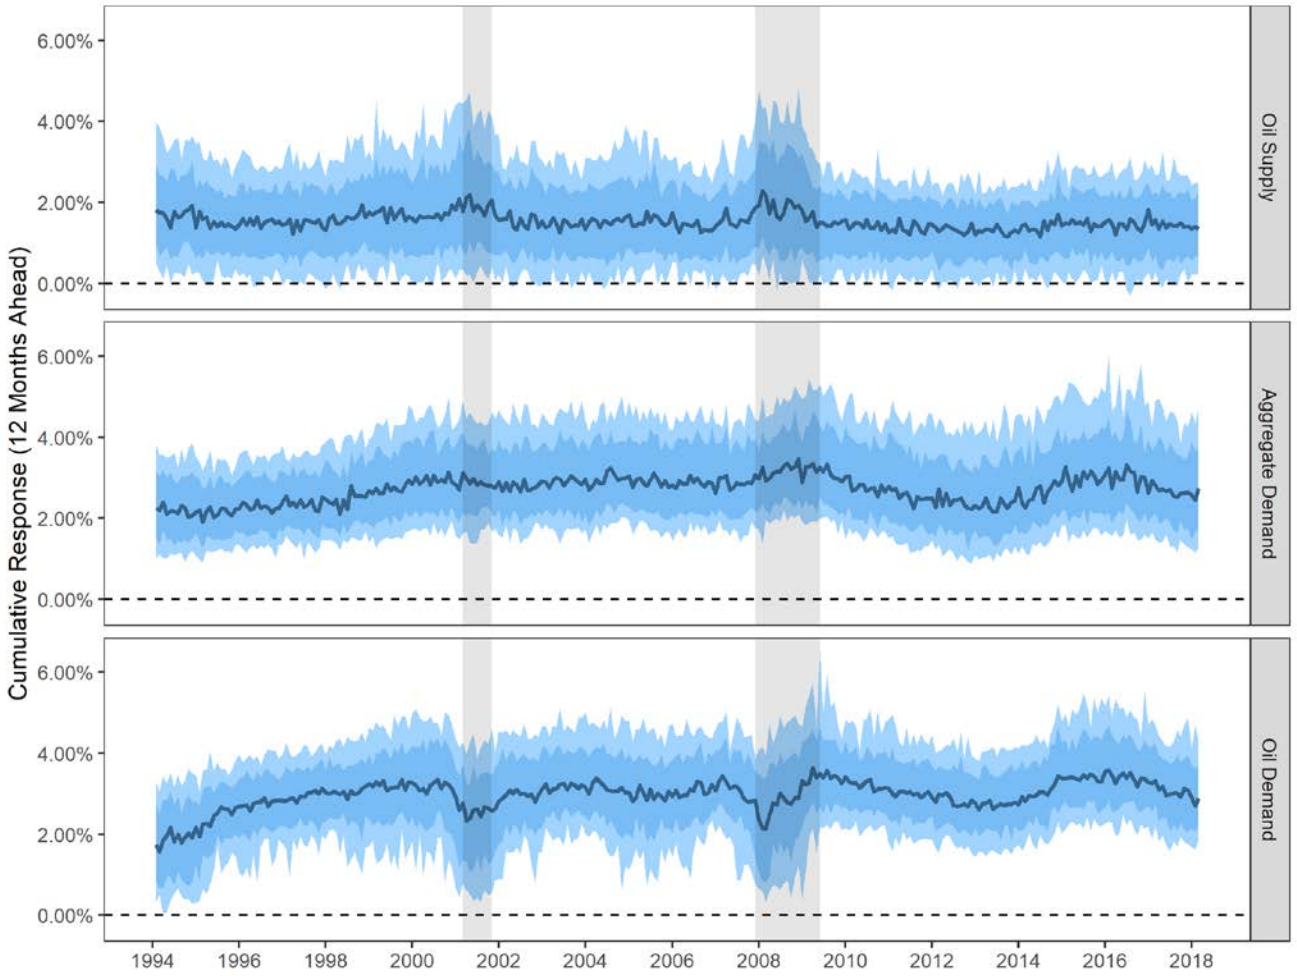

Figure 3.24 Evolution of 12-month cumulative responses of Producer Price Index to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

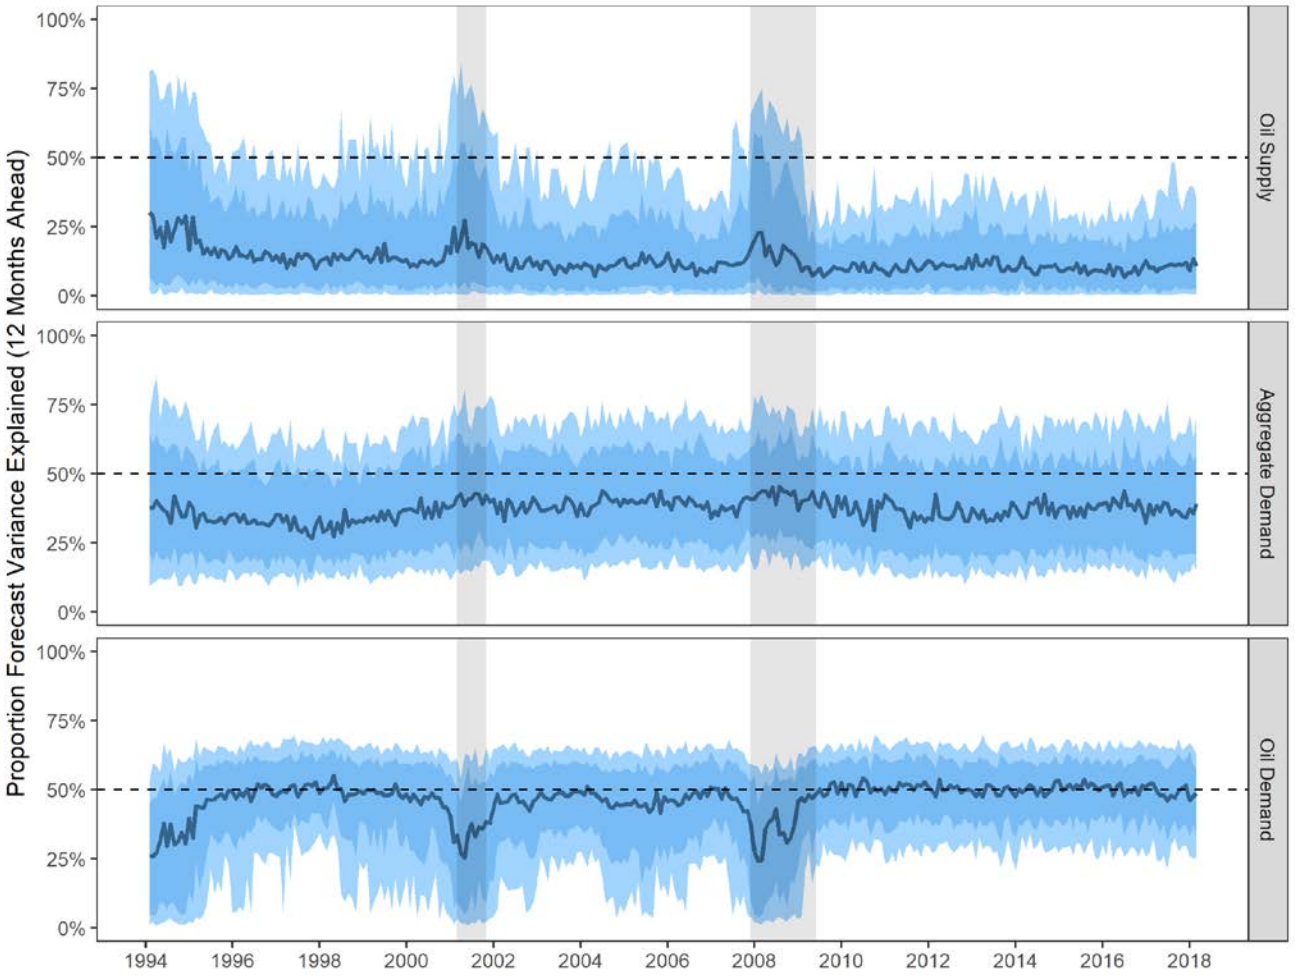

Figure 3.25 Evolution of proportions of the 12-month forecast error variance for Producer Price Index explained by shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

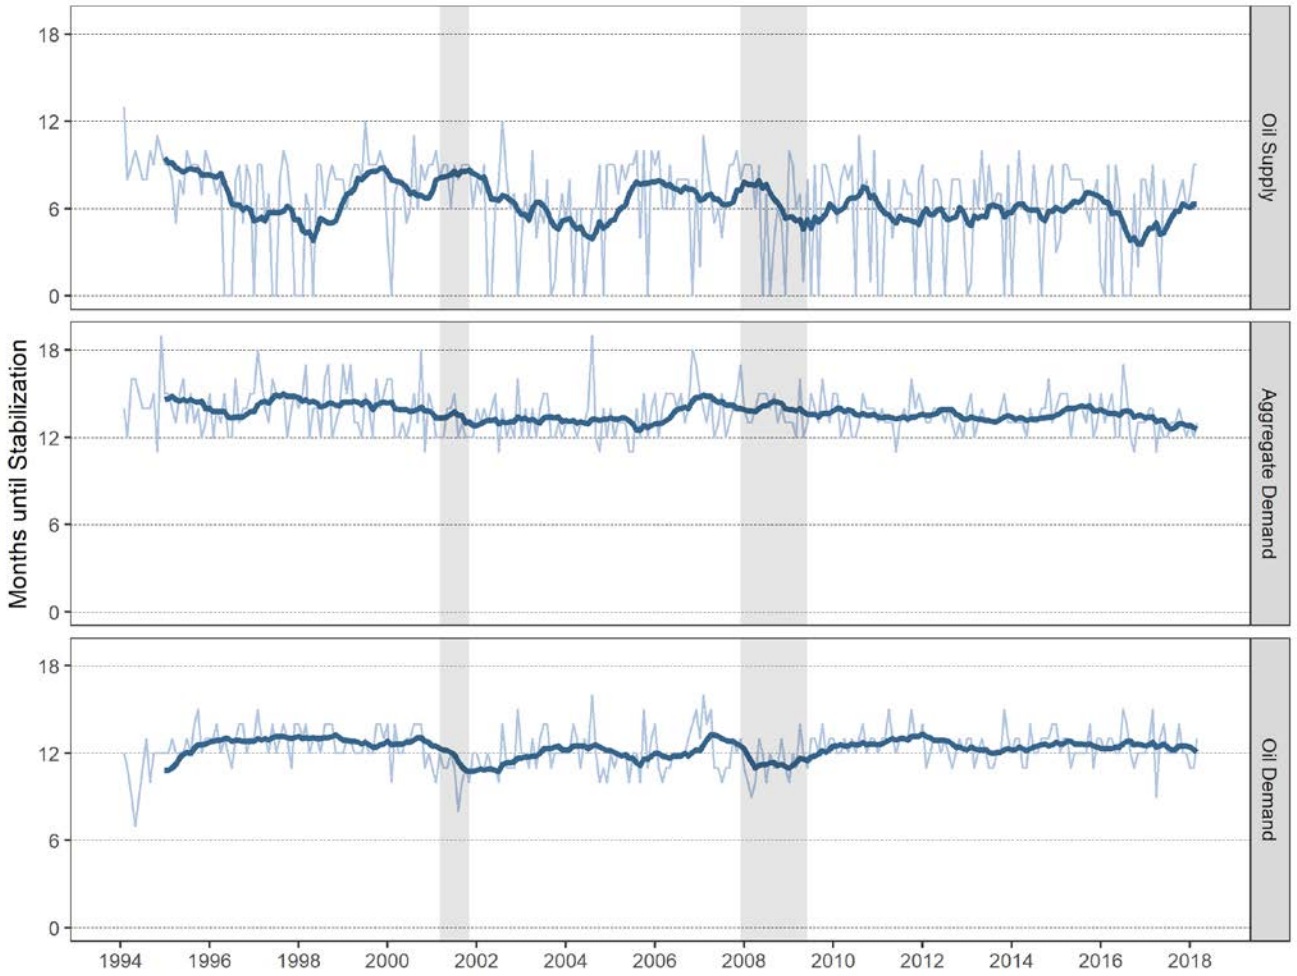

Figure 3.26 Number of months until stabilization of Producer Price Index following shocks in oil supply, aggregate demand, and oil demand with 12-month arithmetic moving average

Table 4.1 Proportion of Graphs with Structural Hamming Distance per Total Relationship (SHDTR) Value of Zero

Data-Generating Process	Graphs Evaluated	Proportion Zero SHDTR
2V1L	2,944,000	0.333
V2L	47,104,000	0.109
3V1L	94,208,000	0.062

Table 4.2 Average Structural Hamming Distance per Lag Relationship by Number of Observations

	Number of Observations				
Data-Generating Process	50	100	250	500	1,000
2V1L	0.314	0.264	0.221	0.204	0.187
2V2L	0.342	0.305	0.271	0.254	0.242
3V1L	0.341	0.294	0.252	0.232	0.217

Table 4.3 Average Structural Hamming Distance per Lag Relationship by the Number of Lags in the PC Algorithm Search

	Search Lags		
Data-Generating Process	Insufficient	Exact	Extra
2V1L		0.140	0.271
2V2L	2.993	0.192	0.283
3V1L		0.181	0.296

Table 4.4 Average Structural Hamming Distance per Contemporaneous Relationship by Number of Observations

	Number of Observations				
Data-Generating Process	50	100	250	500	1,000
2V1L	0.325	0.293	0.255	0.227	0.203
2V2L	0.379	0.358	0.334	0.317	0.301
3V1L	0.382	0.359	0.330	0.310	0.292

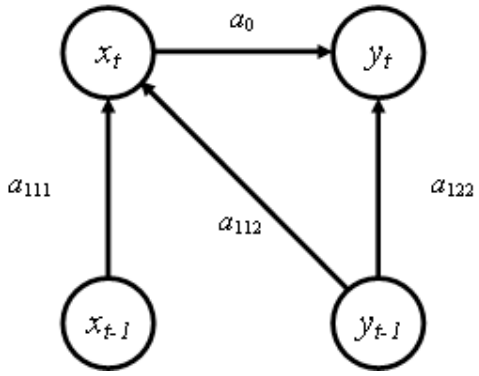

Figure 4.1 Directed Acyclic Graph representation of the data-generating process defined in equation (4.3)

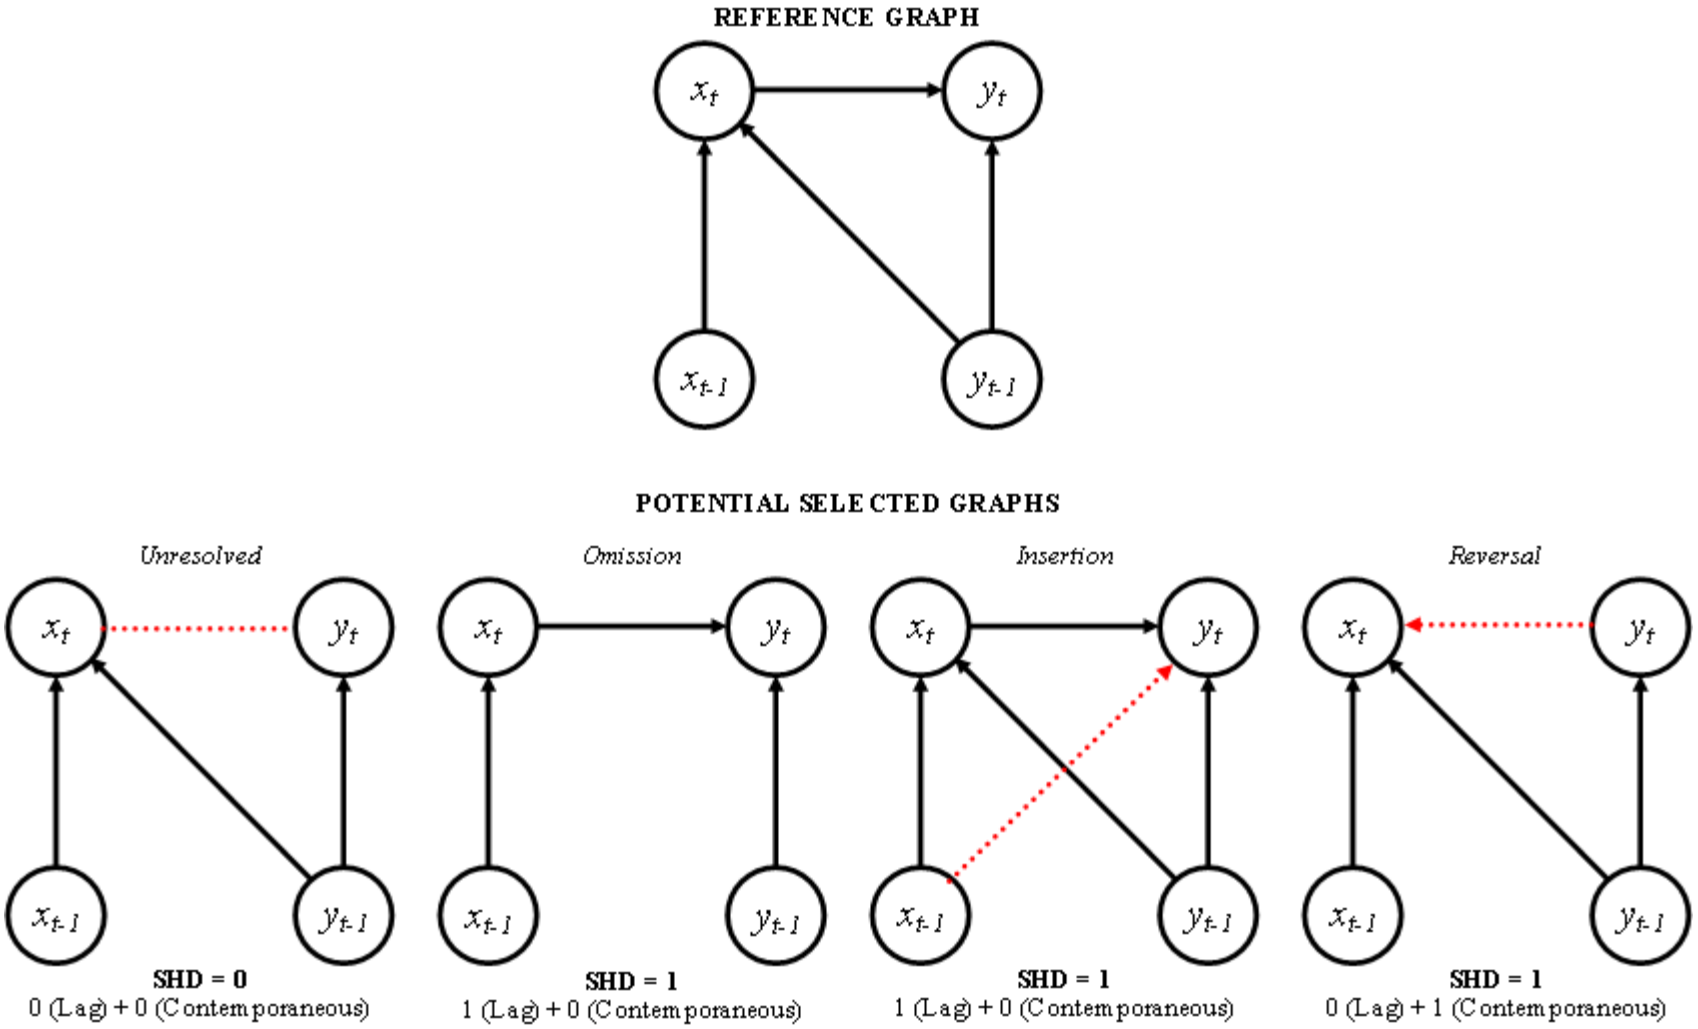

Figure 4.2 Potential errors of the PC Algorithm's selected graph for the data-generating process in equation (4.3)

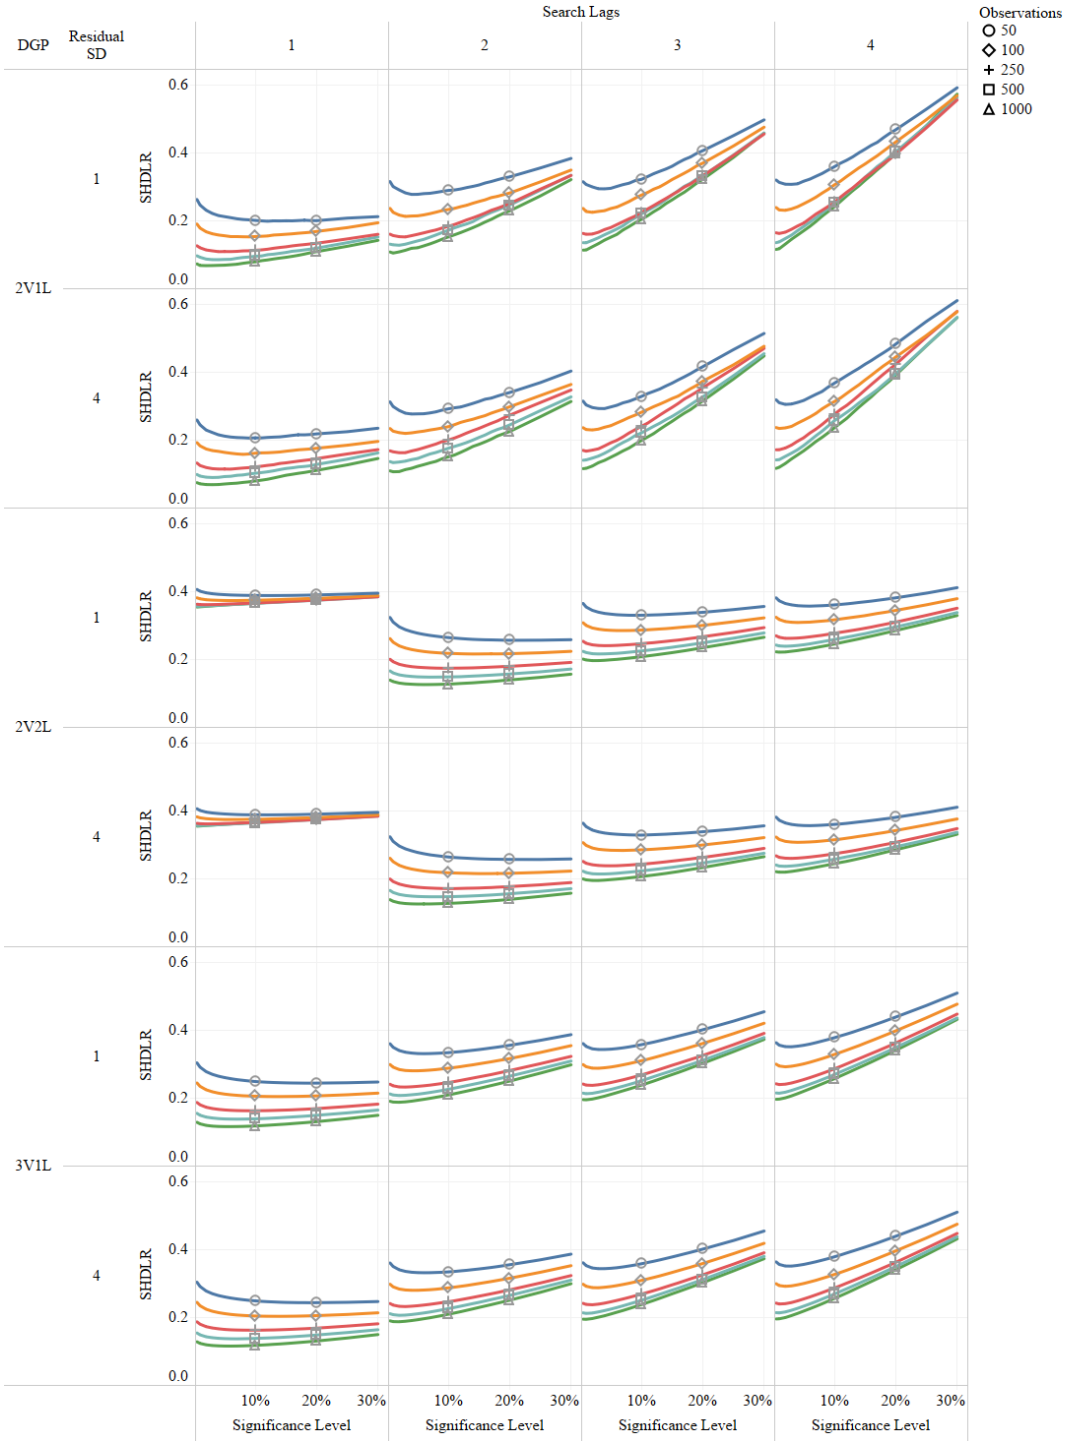

Figure 4.3 Overview of all simulation results for the lag portion of the data-generating process. The values of the Structural Hamming Distance per lagged relationship are averaged, aggregated over data-generating process, the standard deviation of the residuals, the number of lags included in the PC Algorithm, the number of observations in the time series, and the significance level used.

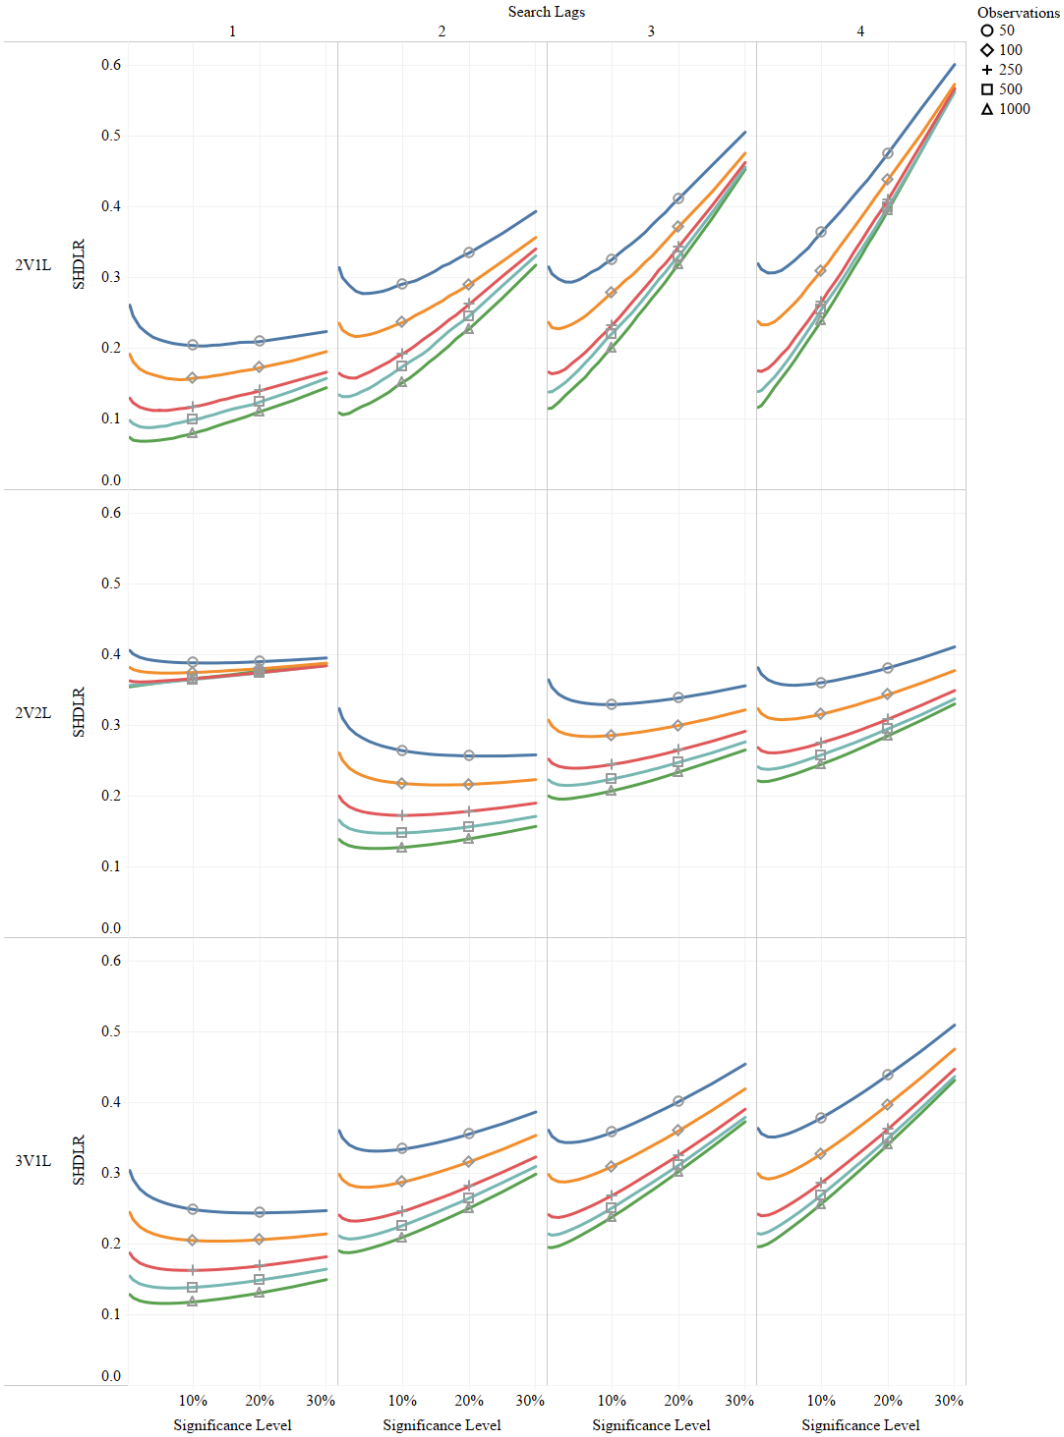

Figure 4.4 Overview of all simulation results for the lag portion of the data-generating process. The values of the Structural Hamming Distance per lagged relationship are averaged, aggregated over data-generating process, the number of lags included in the PC Algorithm, the number of observations in the time series, and the significance level used.

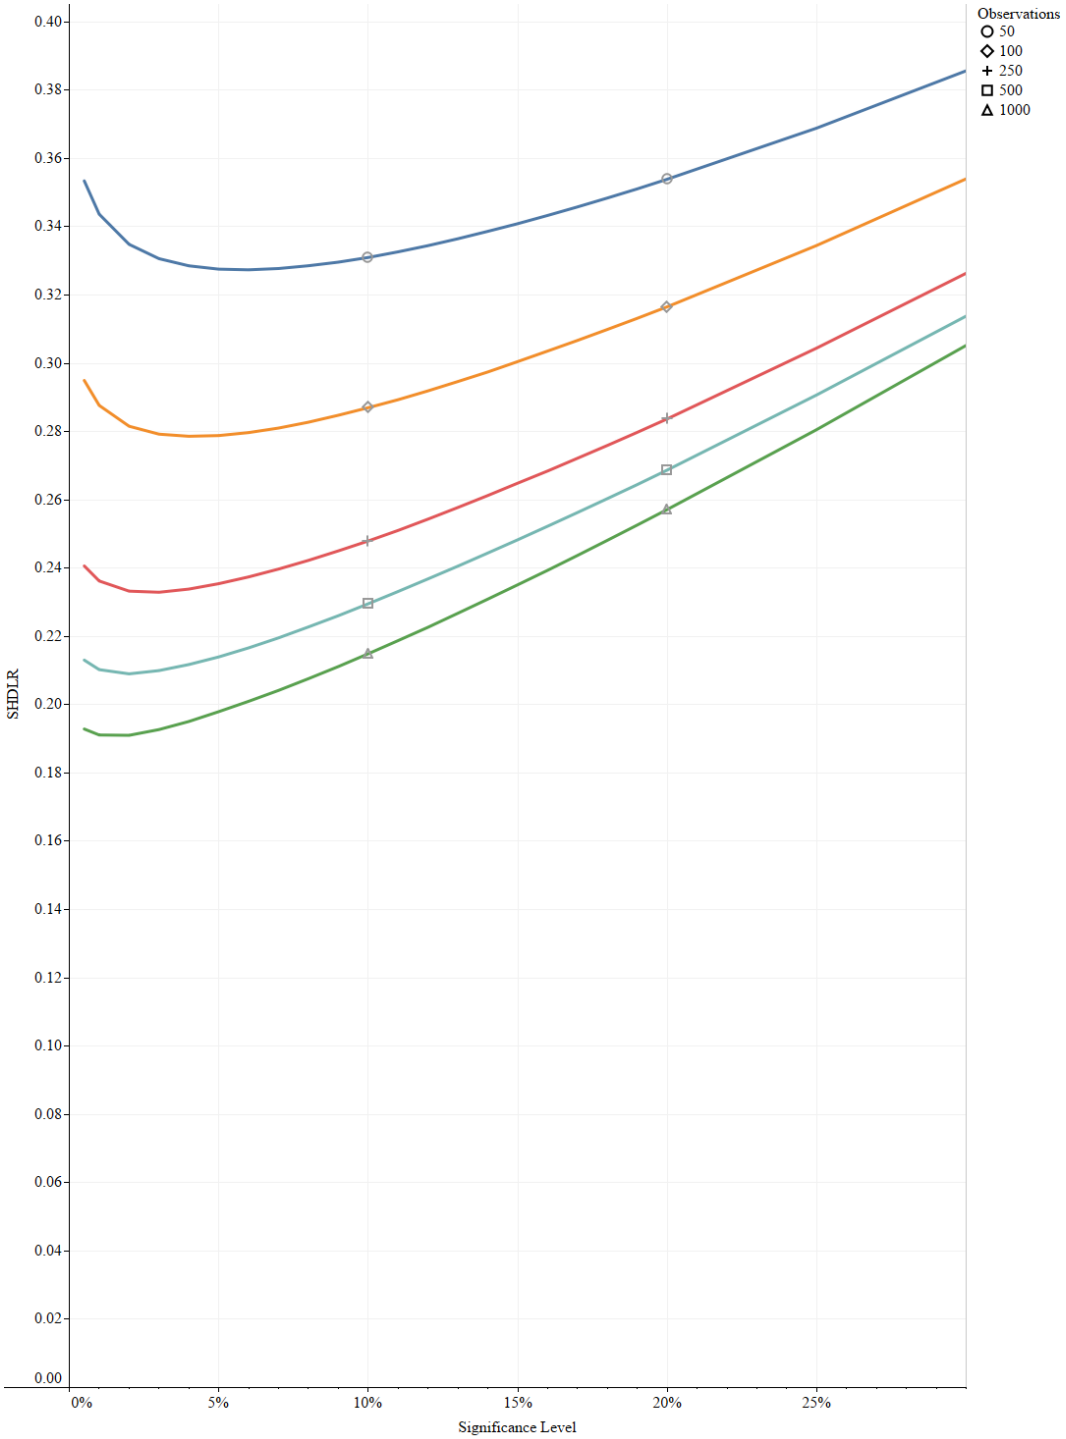

Figure 4.5 Average Structural Hamming Distance per lagged relationship for the lag portion of the data-generating process by significance level

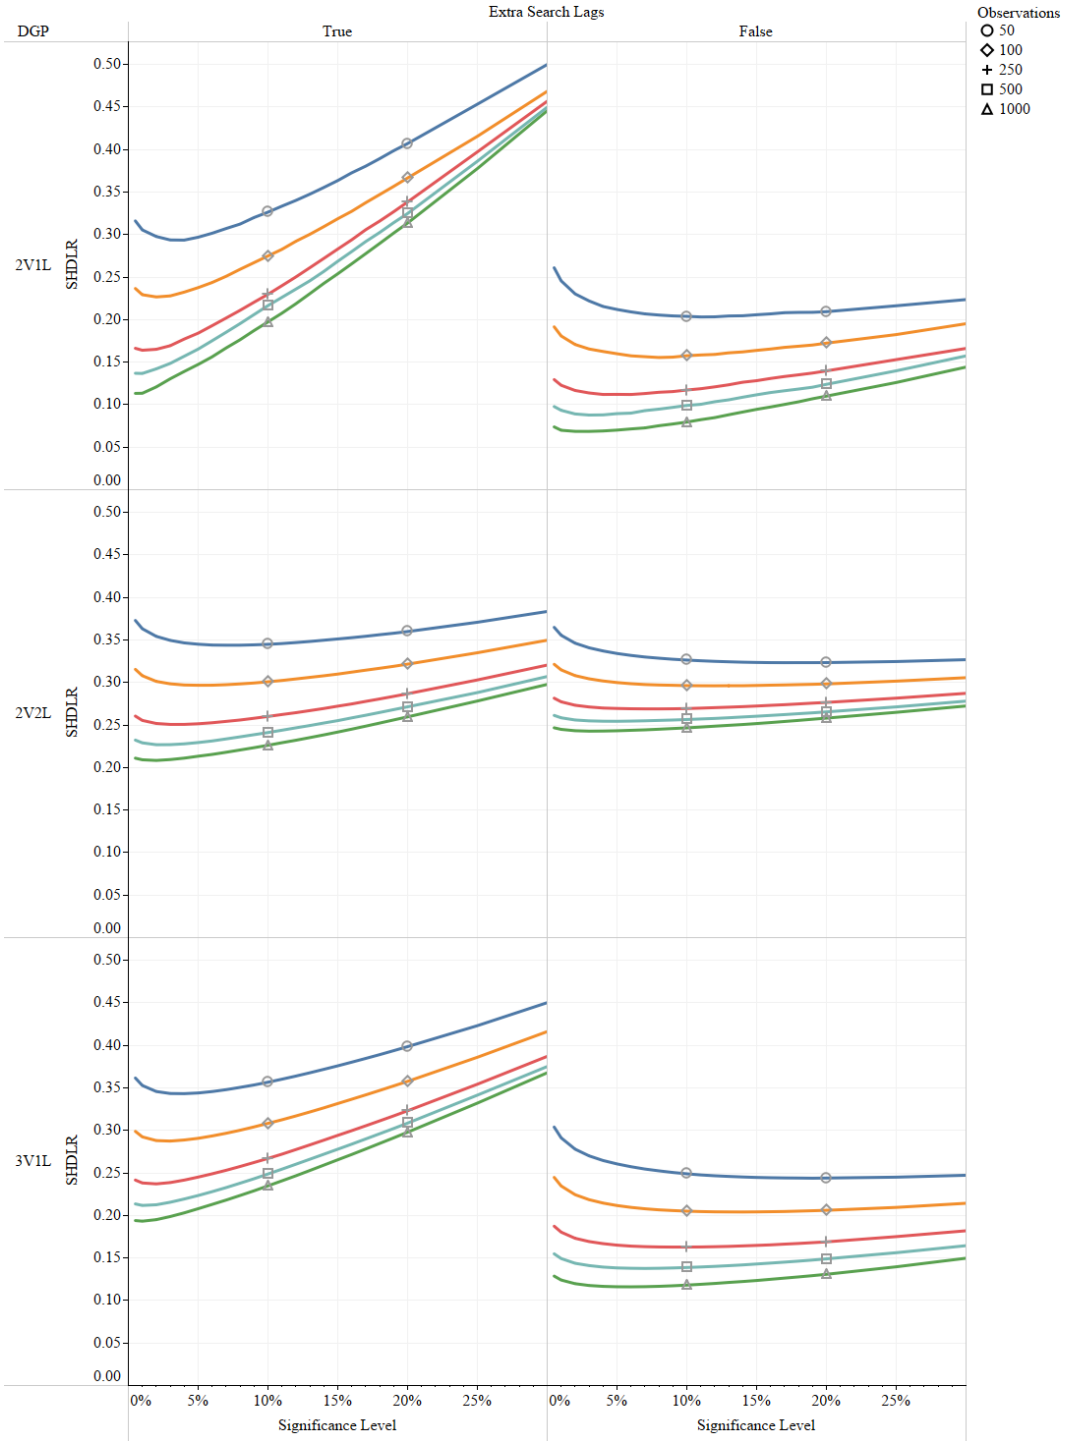

Figure 4.6 Average Structural Hamming Distance per lagged relationship by lags included in the PC Algorithm search. When extra lags is true, the number of lags used in the search exceeds the true number of lags in the data-generating process.

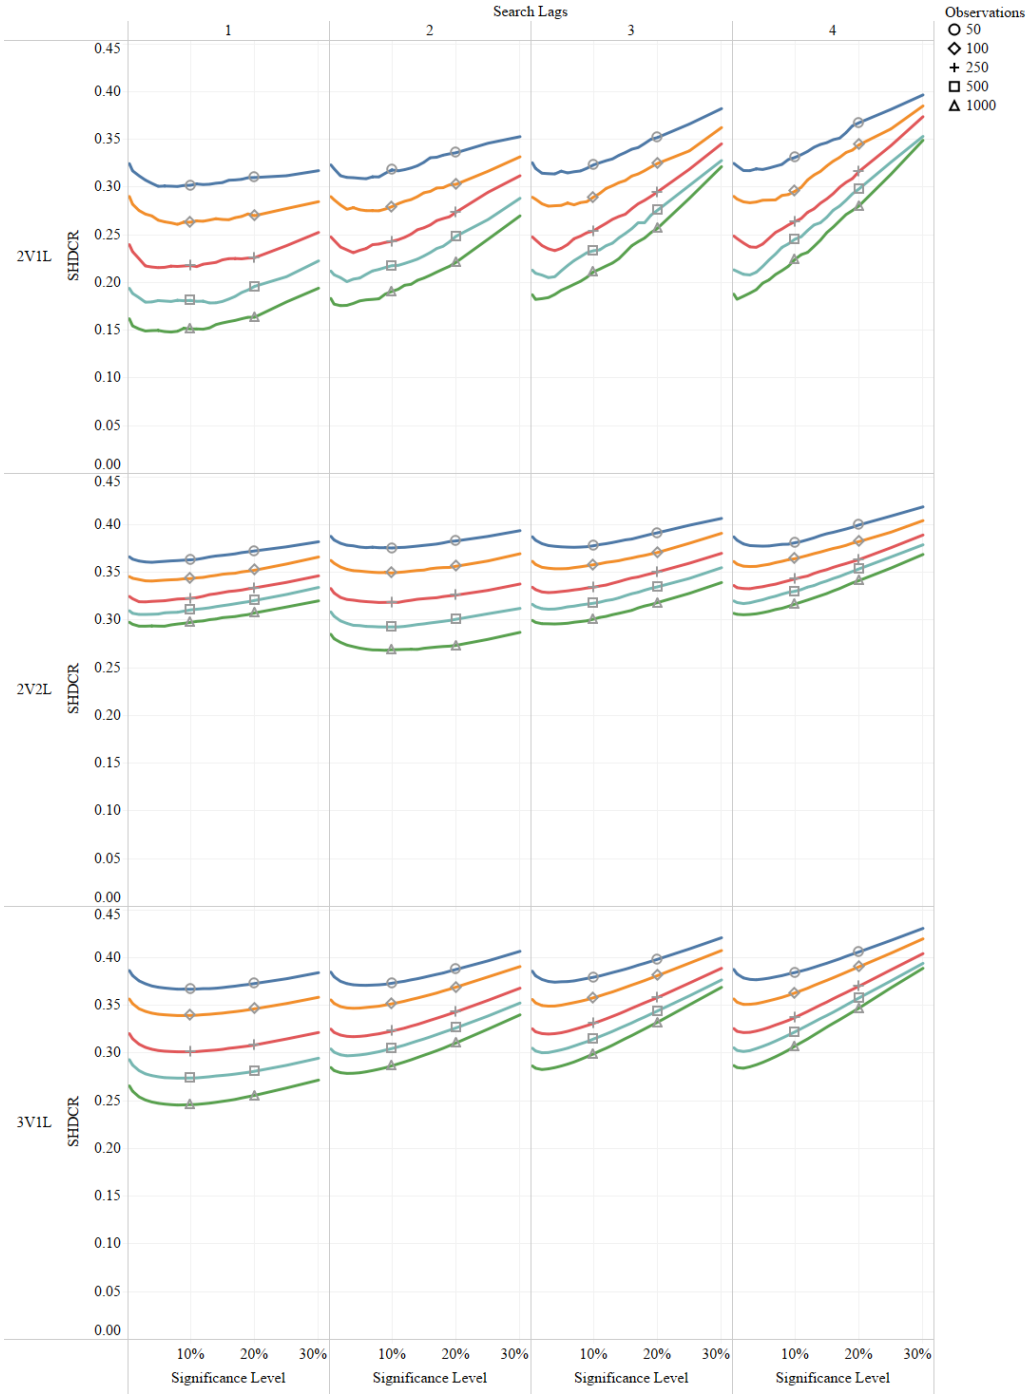

Figure 4.7 Overview of all simulation results for the contemporaneous portion of the data-generating process. The values of the Structural Hamming Distance per contemporaneous relationship are averaged, aggregated over data-generating process, the number of lags included in the PC Algorithm, the number of observations in the time series, and the significance level used.

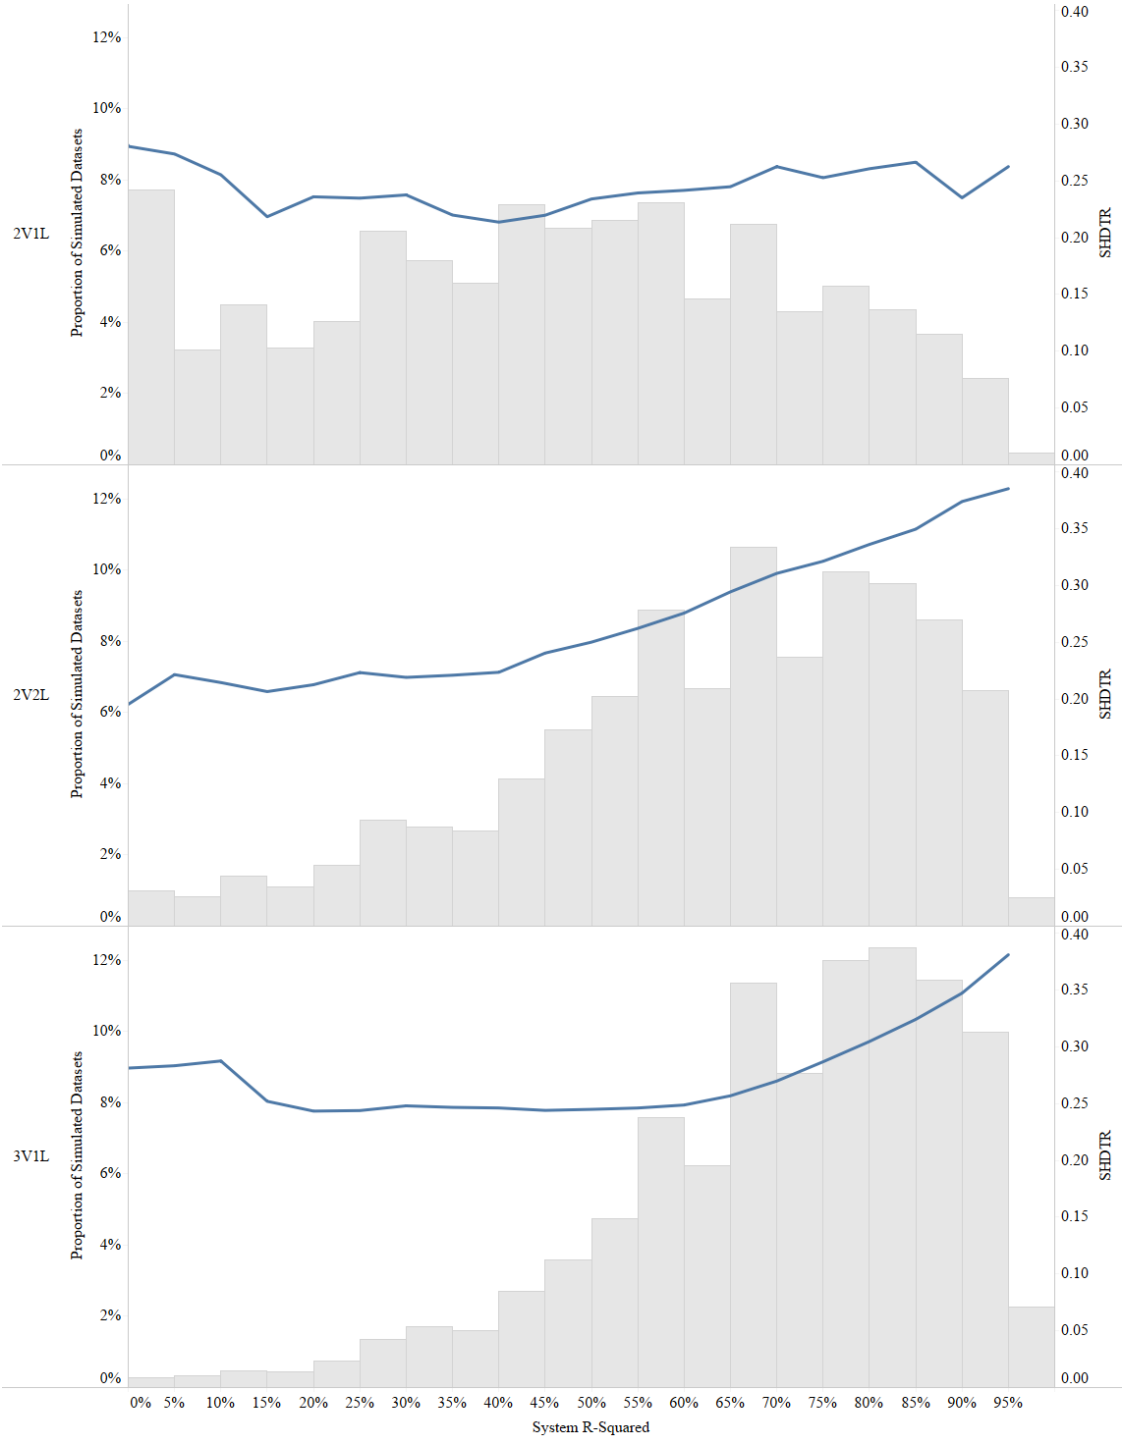

Figure 4.8 Structural Hamming Distance per total relationship for each data-generating process by system R-squared value. The system R-squared values are aggregated into 20 bins five percentage points wide. A histogram with the proportion of all datasets generated falling in each R-squared bin is shown.

APPENDIX B

ADDITIONAL STRUCTURAL IMPULSE RESPONSE FUNCTION PLOTS

For each variable, the structural impulse response functions (SIRFs) are plotted at four selected periods for shocks in oil supply, aggregate demand, and oil demand. Each plot also contains 68- and 90-percent highest posterior densities (HPDs). In general, it is difficult to ascertain the degree of time variation using a small number of selected periods, especially when considering the entire credible set provided by the HPDs.

SIRFs for oil prices in response to these shocks for four selected periods are shown in Figure A.1. A visual examination of these SIRFs reveals no major differences in shape or magnitude for the same shocks across different periods. For an oil supply shock, the median oil prices increase is largest for June 2008, but well within the confidence intervals associated with any of the other three periods. Oil prices respond relatively more to shocks in aggregate demand than to oil supply shocks. The median increase in oil prices for a shock in aggregate demand is the largest in June 2008, but again well within the confidence intervals associated with any of the other three periods. An oil demand shock produces the largest median increase in oil prices of the three shocks regardless of the period.

SIRFs for oil production are plotted in Figure A.2. Again, visual examination reveals no major differences in shape or magnitude for the same shock across different periods.

SIRFs for oil storage at the selected periods are shown in Figure A.3. Storage of oil decreases with all three shocks. Time differences are most pronounced in the response to oil demand shocks; oil storage decreases at the median more for 2009 and 2016 than the other two periods. The effects of oil and aggregate demand shocks are relatively larger in magnitude compared to oil supply shocks.

Capacity utilization for refineries responds to the identified shocks during the same period (Figure A.4). As for the previous variables, the general shapes of the SIRFs are similar across the time periods. The median IRF for 2005 shows a distinct increase after six months that is not as distinct in the other periods. Oil demand shocks produce the smallest responses in magnitude.

SIRFs for gasoline retail sales are similar to those for refinery utilization in shape and differences by period (Figure A.5). The effects of aggregate demand shocks are similar in magnitude but opposite in sign to those of the oil supply shock. Shocks in oil demand decrease refiners' sales of gasoline. The negative response can be explained considering the identifying restrictions for this shock. In this shock, as oil demand increases, so does the price of oil. Because the acquisition costs of oil increase for refiners, their output prices also increase, leading to an increase in gasoline prices for consumers, reducing the quantity demanded of gasoline. At the same time, the demand shock decreases real activity in the economy. Together, these serve to decrease gasoline sales to retail outlets. The effects of oil demand shocks are the smallest in magnitude and shortest in duration of the three shocks.

SIRFs are shown by period for the Industrial Production Index in Figure A.6. None of the responses demonstrate much variability with respect to time. Further, none of the shocks produces a contemporaneous response larger in magnitude than 0.2 percent.

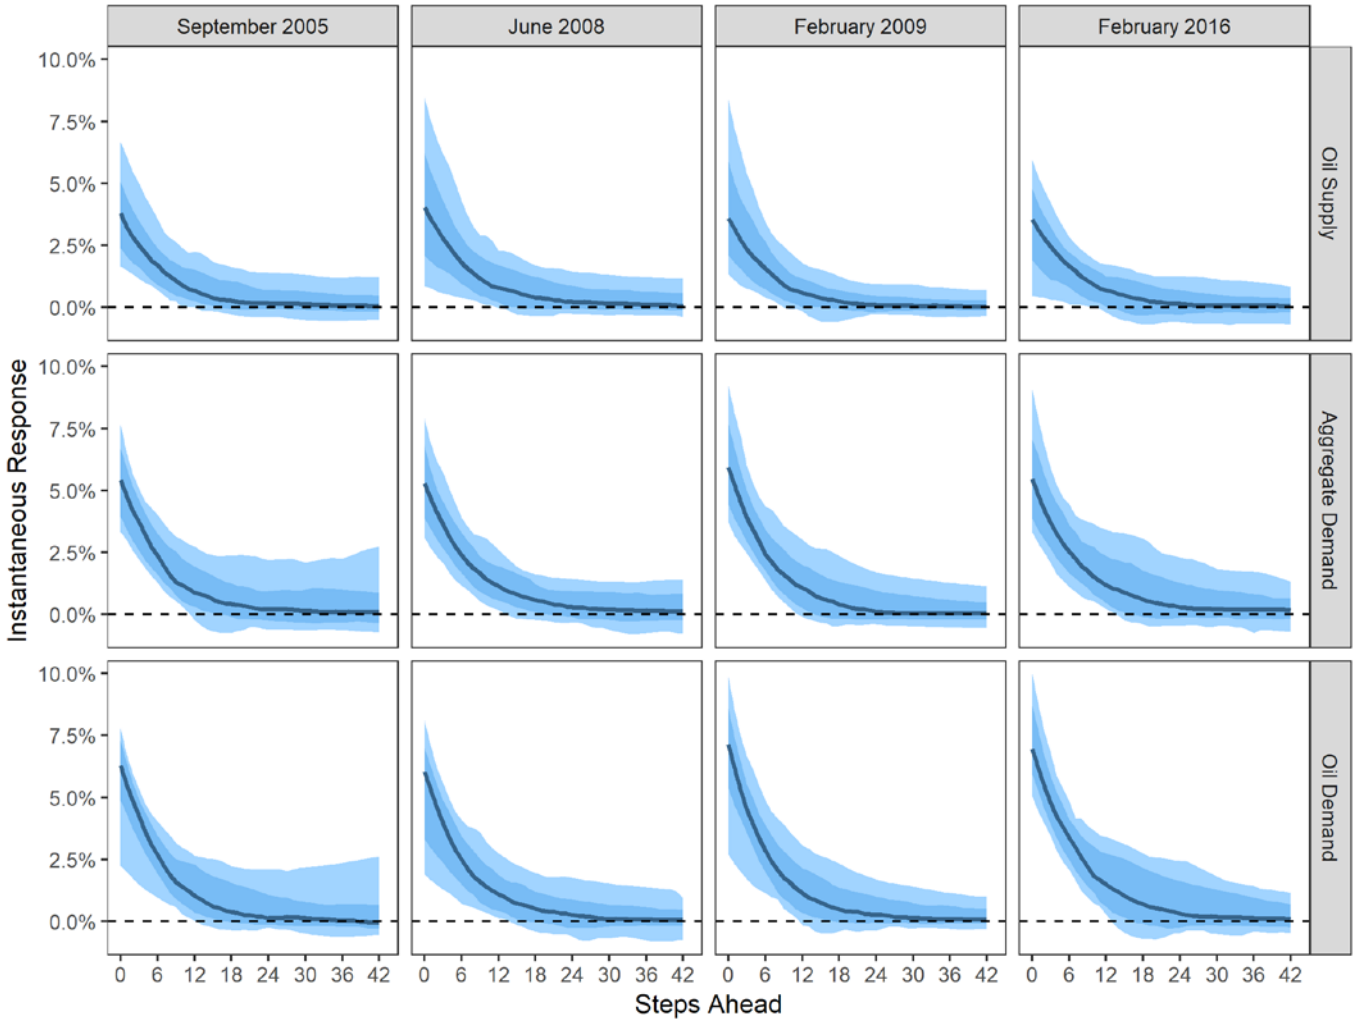

Figure A.1 Instantaneous structural impulse response functions for oil prices to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

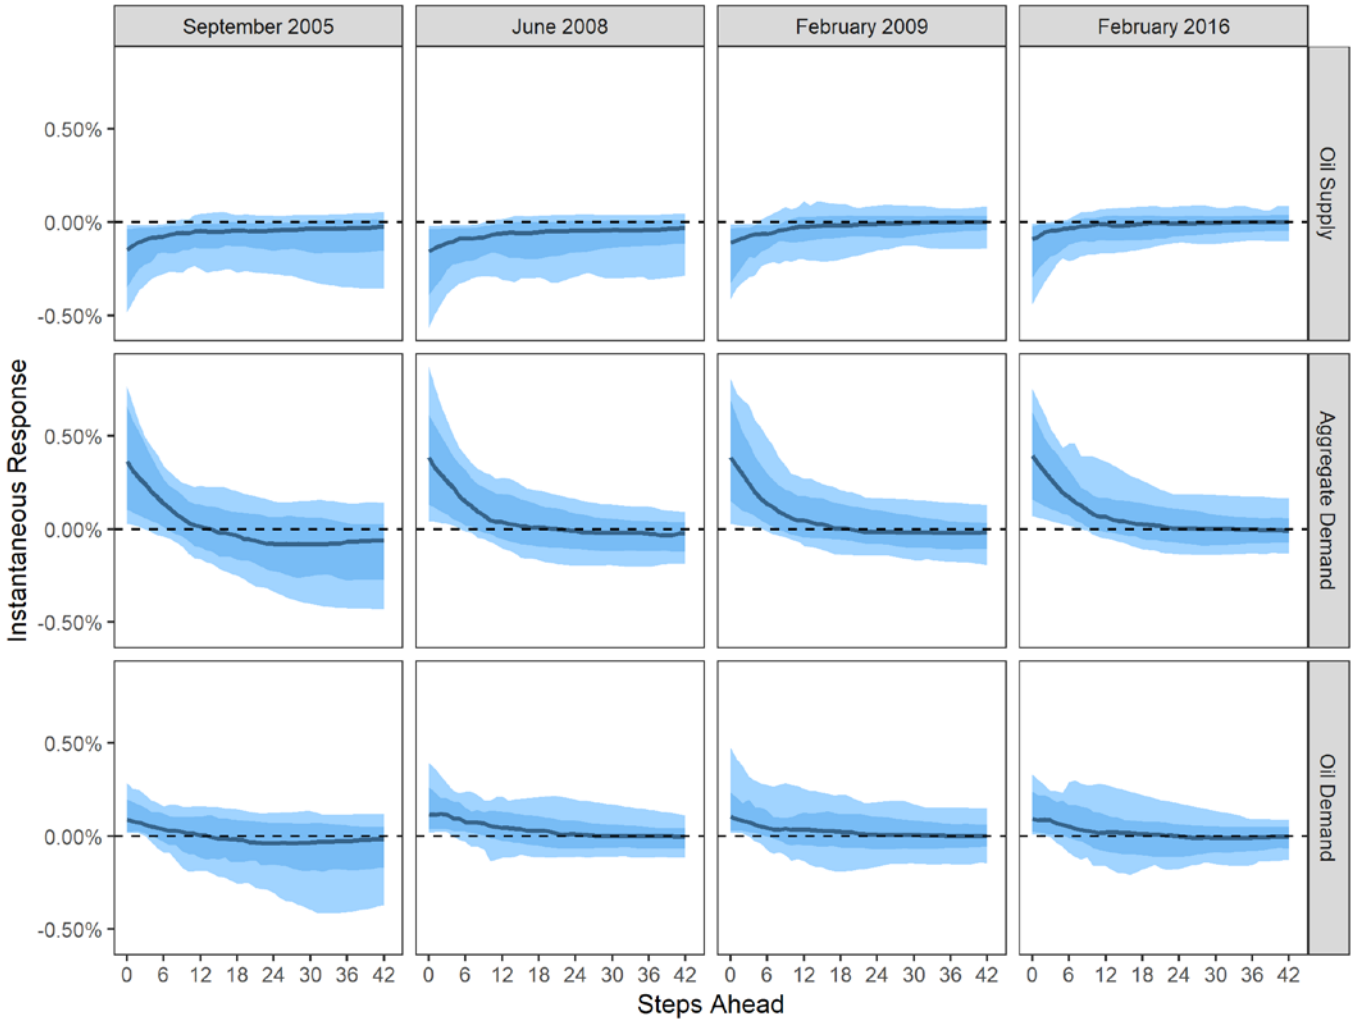

Figure A.2 Instantaneous structural impulse response functions for oil production to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

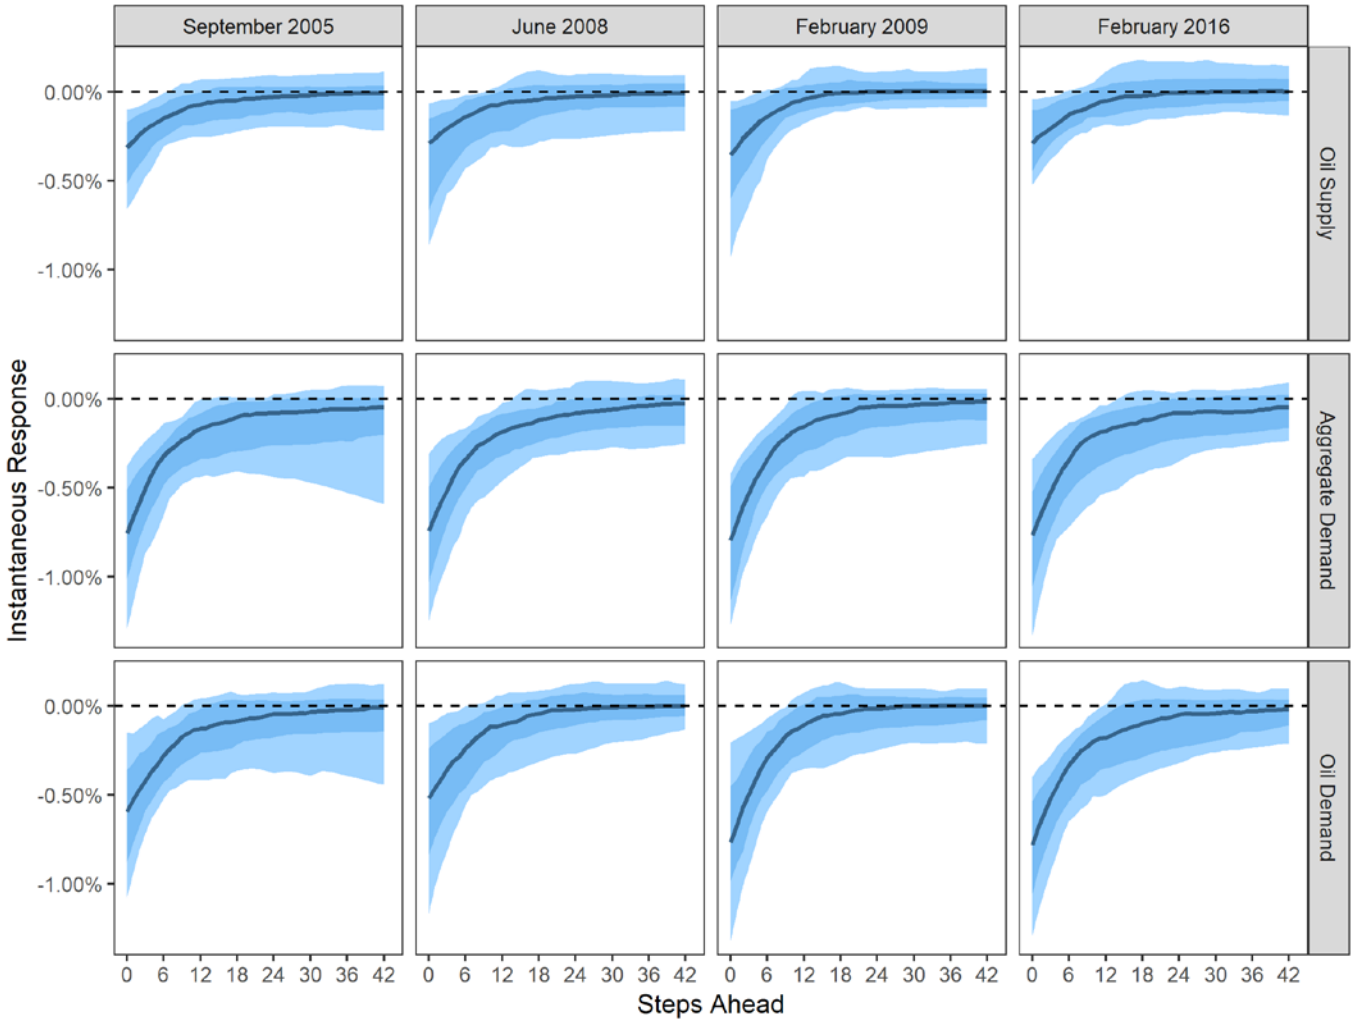

Figure A.3 Instantaneous structural impulse response functions for oil storage to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

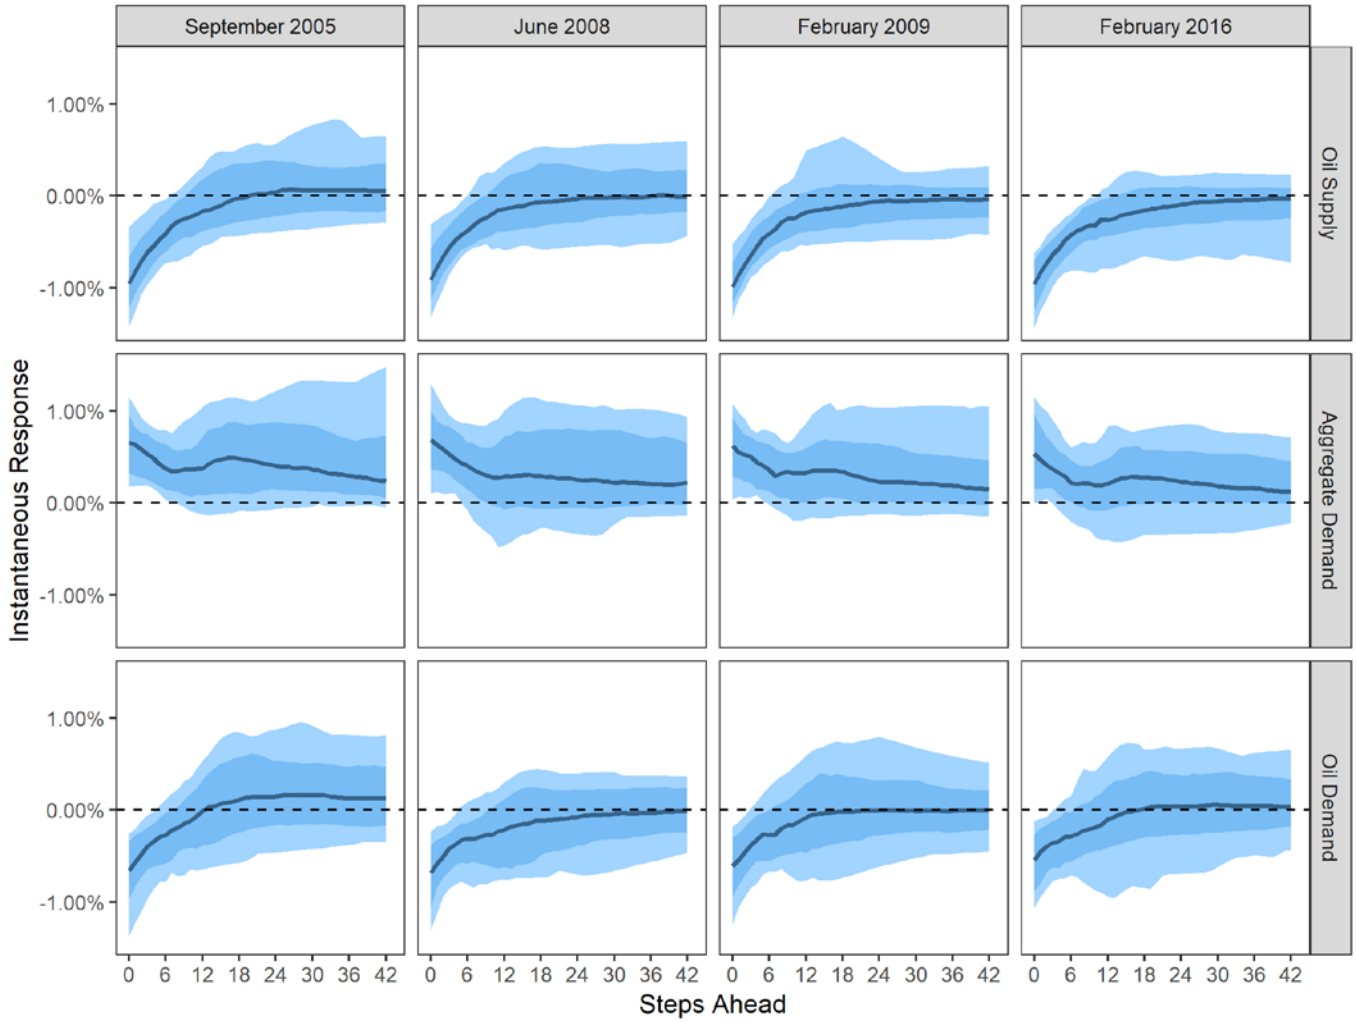

Figure A.4 Instantaneous structural impulse response functions for refinery capacity utilization to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

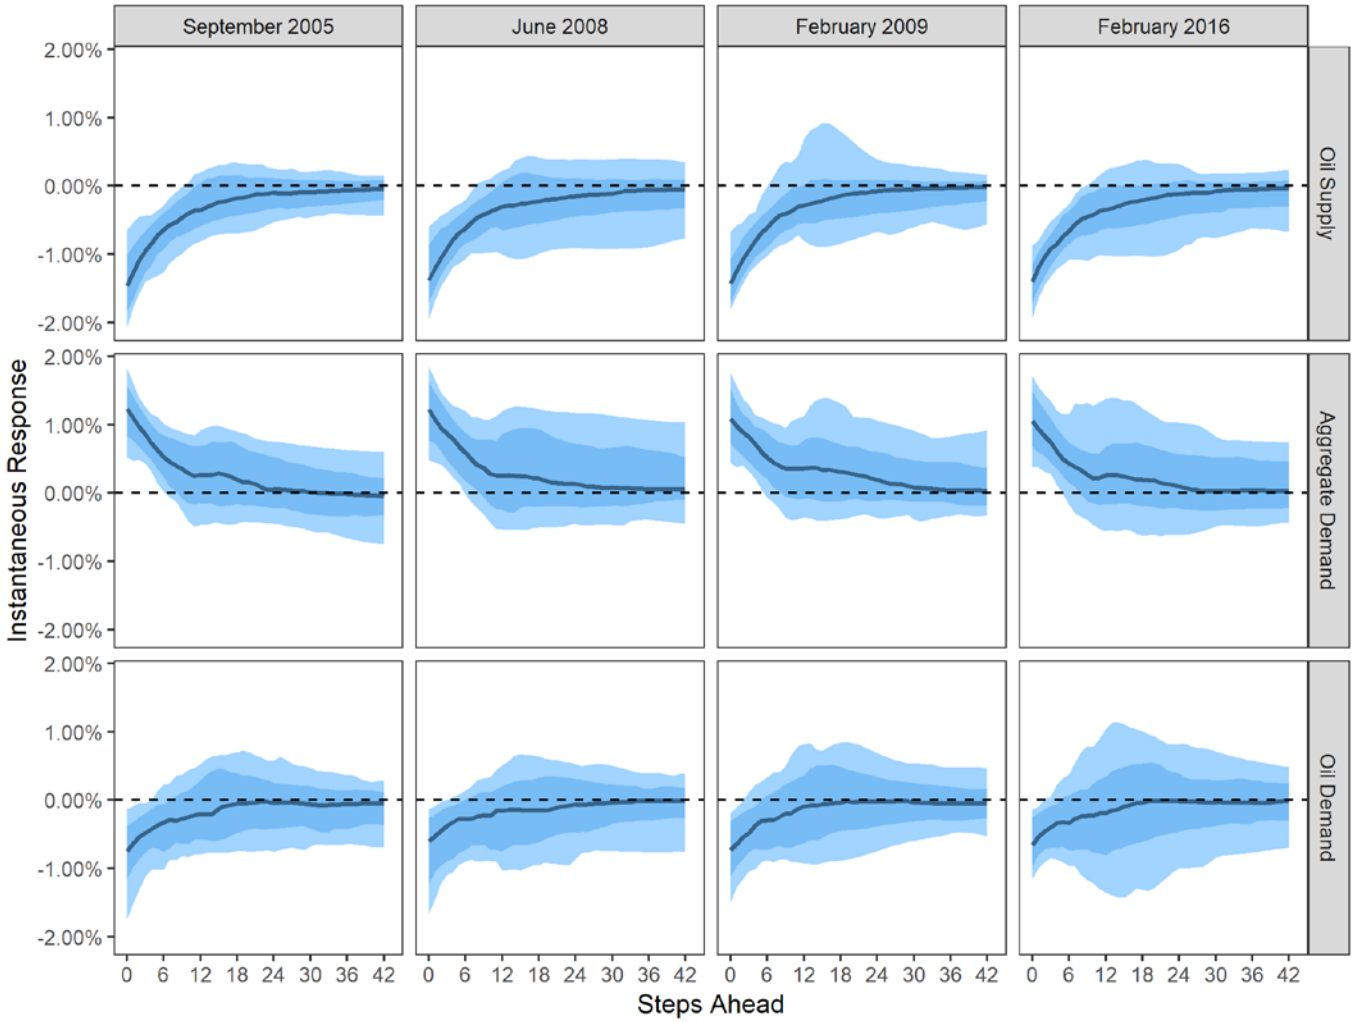

Figure A.5 Instantaneous structural impulse response functions for refiners' retail sales of gasoline to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities

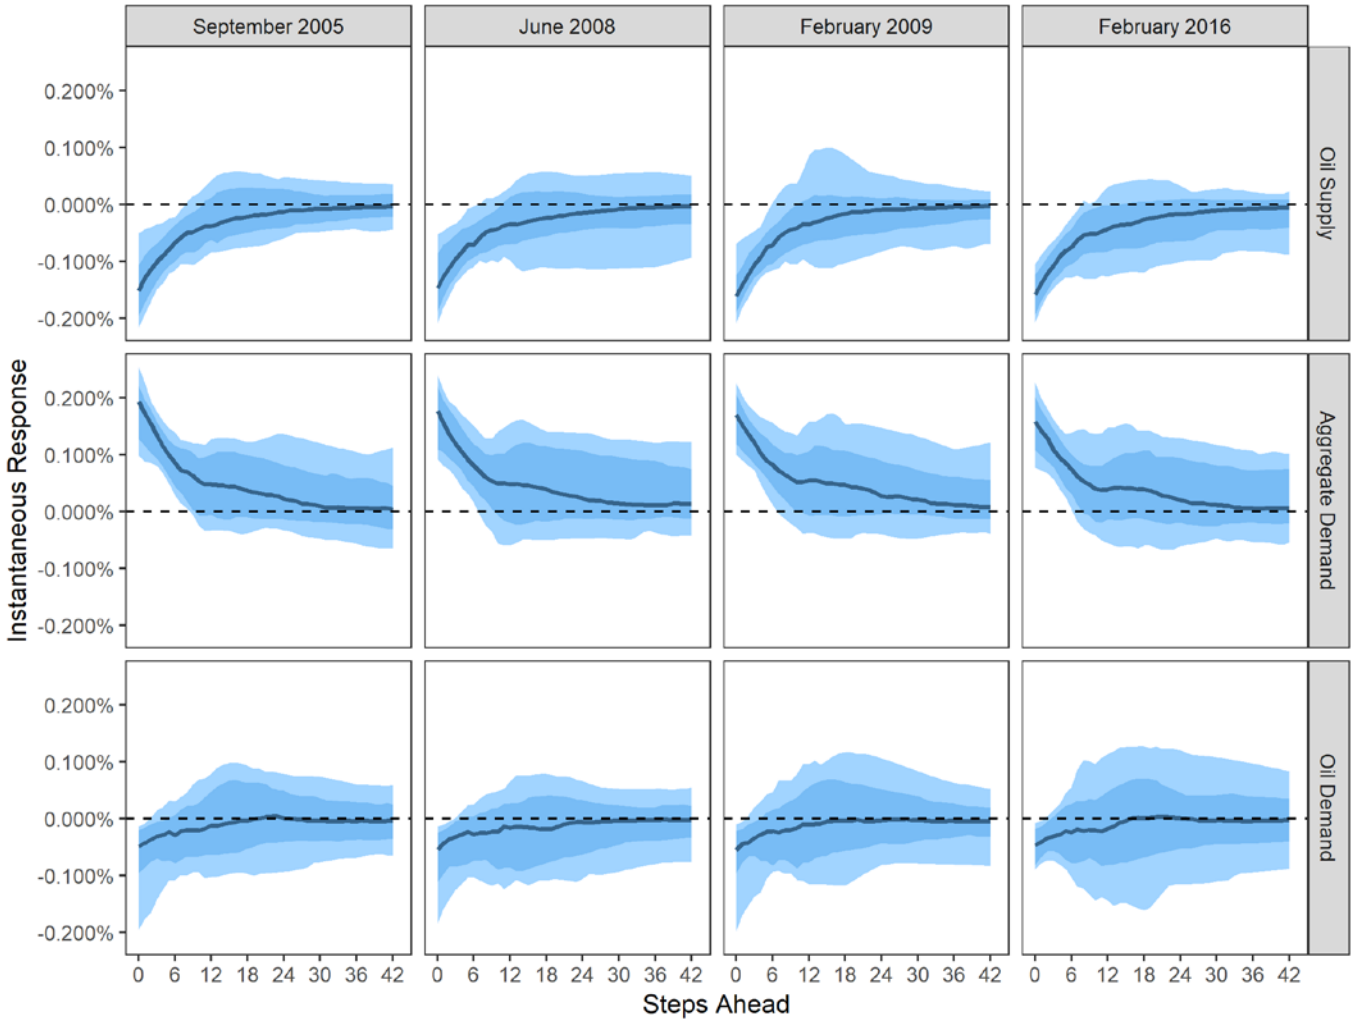

Figure A.6 Instantaneous structural impulse response functions for the Industrial Production Index to shocks in oil supply, aggregate demand, and oil demand with 68- and 90-percent highest posterior densities