



Munich Personal RePEc Archive

Electric Cars and Oil Prices

Azar, Jose

Princeton University

6 August 2009

Online at <https://mpra.ub.uni-muenchen.de/15538/>
MPRA Paper No. 15538, posted 10 Aug 2009 09:31 UTC

Electric Cars and Oil Prices*

José Azar

Princeton University

jazar@princeton.edu

First Draft: January 2009

This Version: August 6, 2009

Abstract

This paper studies the joint dynamics of oil prices and interest in electric cars, measured as the volume of Google searches for related phrases. Not surprisingly, I find that oil price shocks predict increases in Google searches for electric cars. Much more surprisingly, I also find that an increase in Google searches predicts declines in oil prices. The high level of public interest in electric cars between April and August of 2008 can explain approximately half of the decline in oil prices during the second half of 2008. These findings are significant because they show that oil markets respond to developments related to alternative technologies. I investigate several hypotheses explaining these results.

*I am grateful to Chris Sims for invaluable feedback. I thank Pooyan Amir Ahmadi, Francesco Bianchi, Markus Brunnermeier, Jan De Loecker, Penny Goldberg, Andriy Norets, Hal Varian, Mark Watson, and seminar participants at Princeton University for helpful comments and suggestions.

1 Introduction

This paper is an empirical study of the joint dynamics of oil prices and interest in electric cars. Typically, studies of the determinants of oil prices have focused on macroeconomic conditions and supply disruptions rather than developments related to alternatives to oil.¹ I argue that these developments, technological as well as political, are relevant to understanding oil price fluctuations. I also show evidence that interest in alternatives is responsive to changes in the price of oil.

To learn about the joint dynamics of oil prices and the level of interest in electric cars, I use a Bayesian Vector Autoregression (BVAR) with these two variables treated as endogenous. I measure interest in electric cars using data –made publicly available by Google– on the volume of internet searches for phrases related to electric cars, such as ‘electric car’, ‘electric cars’, ‘electric vehicles’, and so on.²

The main findings are as follows. Oil prices respond negatively to shocks to the level of interest in electric cars. The response is economically significant. After a typical shock to interest in electric cars, oil prices decline –gradually– for around six months, reaching a minimum estimated decline of around 5%. This magnitude is of the same order as the typical oil shock in the system. A shock to oil prices, in turn, raises the level of interest in electric cars for around four months. A simulation using only the estimated shocks to interest in electric cars, and no oil shocks, shows that approximately half of the decline in oil prices during the second half of 2008 can be explained by an increase in the level of public interest in electric cars.

The idea that the availability of alternative technologies can play a role in the determination of oil prices is not new. In the second half of the nineteenth century, the most

¹For recent surveys on the economics of oil prices, see Kilian (2008) and Hamilton (2009b).

²The phrases are selected based on the top searches related to electric cars according to Google Insights, as described in more detail in section 2.

important use for oil was illumination using kerosene lamps. The development of electric lighting by Thomas A. Edison in 1879 and of alternating current by George Westinghouse in 1886 provided a substitute for kerosene lighting which was safer, cleaner, and more convenient. According to Williamson and Daum (1959, pages 680-684), the emergence of electric lighting and other alternatives to petroleum products contributed to the stability of oil prices in the last decades of the nineteenth century, despite this being a period of rapid economic expansion. As shown in Figure 1, prices after 1879 were considerably lower and less volatile than in the previous decades.³

Some argue that a similar move from petroleum to electricity is currently under development in the transportation sector.⁴ In the past two years, many countries and regions—including Israel, Denmark, Australia, California, Hawaii and Ontario—have partnered with a startup company called Better Place to provide infrastructure and incentives for the adoption of electric cars. The Chinese government has also publicized its intentions to invest in infrastructure for electric cars and promote them through incentives. At the same time, virtually all of the major car companies have announced their interest in mass-producing electric vehicles. If these high expectations for electric cars were to be realized, there would be a drastic reduction in the long-run demand for oil.

The endogeneity of oil prices with respect to interest in alternatives is problematic for policymakers. The drop in prices after an increase in interest may make it politically more difficult to support the adoption of alternatives to oil. It has the double effect of shifting the public's attention to other topics and making alternatives less competitive.

This paper provides evidence that supports the hypothesis that public interest in energy alternatives is responsive to oil prices. In addition, it shows that oil prices may respond to

³Dvir and Rogoff (2009) show that there was a statistical break in the volatility and persistence of oil prices at some point around 1877. They, alternatively, argue that these changes occurred because Standard Oil, through the installation of oil pipelines, gained enough market power to break the railroads' monopoly on transportation.

⁴See, for example, Better Place (2008), Deutsche Bank (2008), and Becker (2009).

increases in public interest, thus slowing down the political momentum in favor of alternatives. The price of oil is a crucial ingredient in any cost-benefit calculation regarding alternatives, and its endogeneity should be taken into consideration when evaluating energy policy.

2 Data

I measure interest in electric cars using the volume of Google searches for phrases related to electric cars. To choose relevant keywords, I use the top searches related to electric cars according to Google Insights, which is a tool that, given a phrase, provides a ranking of related keywords that people often query on Google before and after that phrase. The top searches in the ‘automotive’ category related to the phrase “electric cars” are shown in Figure 2.⁵ Given a set of phrases, Google Trends then reports a search volume index. The latter is the number of searches for the given the set of phrases relative to the total number of searches at a given moment in time (the index is normalized so that its mean is equal to one). I use the search volume index for the phrase “electric cars” and for the search terms reported in Figure 2 as a measure of public interest in electric cars.

To illustrate what kinds of behavior the Google Trends data can capture, the volume indexes for searches for “ice cream”, “pumpkin”, “down goose” and “economics” are shown in Figure 3. The seasonal patterns that the searches reflect are what one would expect: searches for ice cream are high during summer; searches for pumpkin are low during most of the year but rise sharply at the end of every October; searches for “down goose” increase during the winter; and searches for “economics” decrease during periods of summer vacation and winter holidays. Note that although these examples show clear seasonal patterns the electric car search volume data does not seem to be seasonal.

⁵Google classifies searches into 27 categories, using natural language processing methods. One of the categories is ‘automotive’, to which I restrict the phrases to avoid searches for remote-control electric toy cars.

Data on internet search volume has been used to track flu epidemics by Polgreen et al. (2008) and Ginsberg et al. (2009).⁶ This data has also recently been used by Choi and Varian (2009) to forecast economic series –such as auto sales– before the official data is released.

Figure 4 shows the measure of interest in electric cars described above, together with oil prices.⁷ The data is weekly and goes from the first week of 2004 to July 2009.⁸ The figure shows how events like Hurricane Katrina, the public introduction of prototypes of the Tesla Roadster –a high-end electric sports car– and the simultaneous release of the movie “Who Killed the Electric Car”, as well as press releases from the startup company Better Place, were all followed by substantial increases in interest in electric cars. These were not the only developments that generated search activity for electric cars, but they give an idea of the kind of developments that could be driving public interest.

3 Empirical Strategy

3.1 Modeling the Joint Behavior of Oil Prices and Searches for “Electric Car”

The empirical framework is a Bayesian Vector Autoregression with oil prices and interest in electric cars as the endogenous variables. Thus, I assume that oil prices and interest in electric cars –both in logs– can be described by a structural equation of the form:

$$\Gamma(L)y_t = c + \epsilon_t, \tag{1}$$

⁶This kind of analysis is currently published in real time by Google on its website <http://www.google.org/flutrends/>.

⁷The data for oil prices is the Cushing, Oklahoma WTI Spot Price FOB, in dollars per barrel, which can be downloaded from the website of the Energy Information Administration.

⁸Weekly search volume data measures the intensity of searches from Sunday to Saturday. I define oil price at time t to be the average of the closing prices during the week of searches assigned time t .

where $\Gamma(L)$ is a matrix valued lag polynomial in the nonnegative powers of the lag operator L . The errors ϵ_t are uncorrelated in the time dimension and have a diagonal variance covariance matrix Λ . One can estimate the reduced form vector autoregression

$$\Gamma_0^{-1}\Gamma(L)y_t = \Gamma_0^{-1}c + \Gamma_0^{-1}\epsilon_t, \quad (2)$$

where Γ_0 is the submatrix of $\Gamma(L)$ corresponding to the zero lag. The latter can be rewritten as

$$y_t = d + B(L)y_t + e_t, \quad (3)$$

where $Var(e_t) = \Sigma$. Here, we have defined

$$d = \Gamma_0^{-1}c \quad (4)$$

$$I - B(L) = \Gamma_0^{-1}\Gamma(L) \quad (5)$$

$$e_t = \Gamma_0^{-1}\epsilon_t. \quad (6)$$

This implies that $\Sigma = \Gamma_0^{-1}\Lambda\Gamma_0^{-1'}$.

I use 16 lags, which for weekly data corresponds to around four months. For identification purposes, I use the assumption that the contemporaneous effect of shocks to interest in electric cars on oil prices is zero. The correlation between the estimated shocks is low (.8%), and the main results don't change if one reverses the order of the variables, or if one uses a symmetric square root instead of the Choleski decomposition.

3.2 Prior

I use a Minnesota prior as described in Sims and Zha (1998), which centers prior beliefs around independent unit root behavior. The parameter values are standard and described

in Appendix A. Qualitatively, the results are not dependent on the prior parameterization, as is shown in Appendix A. The magnitude and level of persistence of the response of oil prices to its own shocks and to shocks to interest in electric cars, however, do depend on the parameters.

4 Results

4.1 Impulse Responses

The posterior densities for the impulse responses are shown in Figure 5. In the density plots, the outer bands represent 90% bands, then 80%, 70%, etc., until they reach 10% bands. These error bands are interpreted as the posterior uncertainty for the impulse responses, as discussed in Sims and Zha (1999), not as classical confidence intervals. A shock to oil prices generates a positive response of interest in electric cars, which dies out after around six months. A shock to interest in electric cars produces a negative response of oil prices. After the shock, the oil price declines gradually for around six months, and then slowly reverts to the mean. The magnitude of the response is economically significant: the median response of oil prices to a one standard deviation shock reaches its peak at around negative 5%, which is comparable to the magnitude of a one standard deviation oil shock.

4.2 Simulations with Only Shocks to Interest in Electric Cars

Figure 6 shows the predicted values for both variables when simulating the model without any shocks. This shows the transient generated by the initial conditions. Figure 7 shows the simulation using only the orthogonalized shocks to interest in electric cars, after subtracting the predicted values with no shocks shown in the previous figure. An increase in interest in electric cars in mid-2006 –due in part to the release of the movie “Who killed the electric

car” and the public introduction of the Tesla roadster– can explain most of the decline in oil prices observed in the second half of 2006. Also, a diminished interest in electric cars in late 2007 and early 2008 can explain some of the increase in prices during that period. The high level of interest in electric cars between April and August of 2008 can explain a part of the subsequent crash in prices. Finally, some of the recent increase in prices can be explained by a low level of interest in electric cars at the end 2008 and the first half of 2009.

Most of the increase in interest in electric cars between April and August of 2008 after the price of oil crossed 100USD is assigned to exogenous shocks to interest, rather than to a response to rising oil prices. However, the response of the level of interest in electric cars to oil price shocks could be non-linear. For example, the response could be higher with oil prices above some threshold. This could happen for at least two reasons. First, there might be a psychological threshold at round numbers like 100. Second, it is also possible that there is a value around which electric cars become cost competitive compared to gasoline cars. The value could be around 100USD per barrel, according to calculations by Deutsche Bank (2008), as shown in Table 1.⁹ If there is a threshold, it is likely to decrease over time, as the cost of producing batteries decreases and their life cycle increases.

4.3 Distribution of the Estimated Shocks

Figure 8 shows quantile-quantile plots of the residuals for both equations. The residuals appear to be non-normal –which is confirmed by conducting a Jarque-Bera test. The distributions have fat tails and more positive outliers than negative ones. This raises the concern that a few outliers could be driving the results. By regressing the squared residuals on their lagged values, I also find evidence of serial correlation in the error variances. To address these

⁹According to the estimates reported in the table, electric cars would be cheaper to fuel –including the cost of electricity and battery depreciation– than gasoline cars at gasoline prices of \$3.50 per gallon or more. Using a regression of oil prices on gas prices in the United States with data for 2008 and 2009, the expected value of oil prices given gas prices at \$3.50 is slightly above 100USD per barrel of oil.

concerns, I re-estimated the vector autoregression allowing for regime-switching in the means and variances of the error terms, which allows for a general error distribution –with both fat tails and asymmetry– and for autocorrelation in the error variances. The impulse responses from the regime-switching model look quite similar to the normal model. The magnitude of the effect is smaller, but still statistically and economically significant. The estimation procedure and results are described in detail in Appendix B.

5 Discussion

5.1 Why Would Interest in Electric Cars Affect Oil Prices?

The fact that interest in electric cars responds to changes in oil prices is intuitive. But why would interest in electric cars, or in alternatives in general, precede declines in the price of oil? I will discuss possible mechanisms that could be driving the results. A conclusive analysis, however, is beyond the scope of this paper.

One possibility is that interest in electric cars may be acting as a proxy for technological, business, and policy developments related to electric cars. These can translate into a higher long-run price elasticity of the demand for oil, because possible alternatives become more convenient or cheaper. This higher long-run elasticity would translate into lower oil prices.

A second possible mechanism is that oil producers with market power may react strategically to increased public interest in alternative technologies, in an effort to slow down the substitution of oil by the former. To do this, they could lower oil prices when there is a substantial increase in public interest in alternatives. Conversely, during periods when interest in alternative technologies is low, prices would go up.

Lower oil prices would work against the adoption of alternatives for several reasons. First, they may make it politically harder to pass legislation to facilitate and incentivize the

adoption of alternatives. With lower oil prices, alternatives become less attractive relative to oil in terms of cost leading to a decline in public interest in alternatives. Second, a decline in oil prices reduces the incentives for car companies to make long-term investment decisions related to the manufacturing of fuel efficient cars. Finally, lower oil prices also put financial strain on startup companies related to alternative technologies. For example, Tesla Motors was having financial troubles in October 2008, having to fire workers and postpone the introduction of their second car model. According to a news article in *BusinessWeek*, “[i]f oil prices continue their fall and stay below \$80 a barrel, Tesla, which currently spends nothing on marketing, could lose some of its buzz.”¹⁰

Whichever the underlying mechanism, the sluggishness of the estimated reaction of oil prices after a shock to interest in electric cars requires an explanation. If the mechanism is through financial markets, then this could be understood in the context of delayed reaction to news. In informationally efficient markets, the reaction of oil prices to shocks to interest in electric cars should be instantaneous. This said, the empirical evidence shows that there is underreaction to news at short-run horizons (three months to a year) in financial markets.¹¹ Hong et al. (2000) show that the underreaction is stronger for stocks with a low level of analyst coverage. Thus, if oil analysts are focused on macroeconomic developments and supply disruptions rather than long-run developments of alternative technologies, it is plausible that the reaction to changes in interest in electric cars would be delayed. Hong and Stein (1999), Barberis et al. (1998) and Daniel et al. (1998) propose behavioral models that can account for this behavior. Theories of rational inattention could also be a useful theoretical framework to model delayed reaction to news in financial markets.

If the mechanism is a strategic reaction by oil producers, the same arguments hold. Due to natural limitations in information processing capacity, they could become aware of changes in

¹⁰David Welch, “Electric Carmaker Tesla Downshifts,” *BusinessWeek*, October 23, 2008, Section Autos.

¹¹See, for example Cutler et al. (1990) and Jegadeesh and Titman (1993).

interest in alternatives gradually, and take time to react. Moreover, the speed of the reaction might be strategic as well, since oil producers have an interest in keeping price volatility low.

5.2 Why Did Oil Prices Rise and Fall in 2008?

As shown in the simulations in Figure 7 described in Section 4.2, some of the increase in oil prices in 2007 and the first half of 2008, and a substantial part of the fall in oil prices in the second half of 2008 can be explained by changes in interest in electric cars. This is consistent with the view expressed by Hamilton (2009a) that the rise and fall in oil prices during this period were likely due to changes in the long-run elasticity of the demand for oil. In particular, he argues that the economic downturn by itself is not sufficient to explain the magnitude of the decline in oil prices observed in the second half of 2008. Instead, he attributes the price collapse to an increase in the long-run elasticity of the demand for oil.

It is interesting to compare the behavior of spot and futures prices during this period. Futures prices for contracts with a maturity of several years provide information on the traders' for the oil market many years into the future. Figure 9 shows the prices oil futures contracts with times to maturity of 35 and 72 months, together with the spot price.¹² The rise in prices until mid-July of 2008 was similar for spot and futures prices. The decline of futures prices starting in mid-July of 2008 was also similar to that of spot prices –although spot prices fell more. This suggests that both the rise and decline in oil prices during 2008 were driven primarily by changes in expectations regarding the long-run demand and supply of oil, rather than short-run considerations such as business cycles or capacity constraints.

A way to examine the relative importance of the recession versus developments related to alternatives is to run a three-variable Vector Autoregression adding a measure of economic activity as an endogenous variable. There are no official measures of economic activity at

¹²There are missing values for some days in the Bloomberg data for long-term futures contracts. The 35-month contract had data for most days of the 2004-2009 period; contracts with a maturity of 36 months or more have more missing values.

the weekly frequency, but Aruoba et al. (2008) calculate a filtered daily index of business conditions, published by the Federal Reserve Bank of Philadelphia.

Figure 10 shows the resulting impulse responses. Figure 11 shows the counterfactuals of oil prices with shocks to interest in electric cars only, and with only shocks to the Aruoba-Diebold-Scotti Business Conditions Index. As can be seen from the figures, shocks to interest in electric cars still have a significant negative impact on oil prices, and can explain a large fraction of the decline in the price of oil during the second half of 2008. A shock to business conditions has a positive impact on the oil price. The orthogonalized shocks to business conditions can account for some of the price decline during the second half of 2008. However, the fraction of the price decline that can be explained by shocks to business conditions is substantially smaller than that which can be explained by shocks to interest in electric cars. Also, in the simulation using only shocks to business conditions, oil prices start to decline at the end of 2007, while actual oil prices rose steeply until mid-July of 2008.

6 Summary and Conclusion

This paper shows evidence that oil prices respond negatively to shocks to the level of public interest in electric cars. It also shows that public interest in electric cars responds to shocks to oil prices. Changes in the level of interest in electric cars can explain around half of the price decline in the second half of 2008. These results suggest that oil prices are not driven only by macroeconomic factors and supply risks. Technological, business, and political developments related to alternative technologies may be just as relevant in understanding changes in oil prices.

I have provided only an informal discussion of the underlying mechanisms that could explain the results. A more thorough exploration is a natural avenue for further research. The results are likely to have implications for policy analysis. Clearly, one of the main ingredients

in any cost-benefit analysis regarding alternatives to oil is the price of oil itself. However, the price of oil is endogenous with respect to policy related to alternative technologies. This paper presents evidence that suggests that the magnitude of the effect of policy on prices could be economically significant. This price endogeneity has at least two implications for policy analysis. The first is that it represents a change in the terms of trade favorable to oil importers. The second is that it slows down the adoption of alternatives. This makes the policy less effective in promoting alternatives than what a policymaker that treats the price of oil as exogenous would expect.

Appendices

A Minnesota Prior

A.1 Parameterization for the Reported Results

The parameterization of the prior is standard.¹³ The weights on the sum-of-coefficients dummy λ , the weight on the no-cointegration dummies μ , and the overall tightness parameter on the ‘Minnesota-prior’ dummies π_1 are set to 5, 2, and 3, respectively. The parameter that allows different weights between own lags and coefficients on other variables, π_2 , is omitted, which is the same as setting it to 1. This is required to be able to implement the prior using system-wide dummy observations. The decay parameter π_3 is set to one half. The σ_k parameters, which measure variability of variable k , are set equal to the standard deviations of the differences in each variable. The weight on the variance-covariance matrix dummy w is set to 1. The values assigned are summarized in Table A.1.

A.2 Sensitivity to Different Values for the Hyperparameters

Figures A.1 to A.4 show the response of oil prices to its own shocks and to shocks to interest in electric cars for different values of the hyperparameters. The impulse responses are almost identical for different values of λ , π_1 , and π_3 . Changing the weight μ on the no-cointegration dummies affects the persistence of the response of oil prices to both types of shocks, with higher values implying a higher level of persistence. A value of $\mu = 0$ imply fast mean reversion of oil prices, with a typical shock vanishing completely around 80 weeks after impact.

¹³The notation follows that of Chris Sims’ lecture notes and VAR tools –available at <http://sims.princeton.edu/>– rather than that in Sims and Zha (1998).

B Regime-Switching Vector Autoregression

This section describes the Vector Autoregression model with regime-switching for the error means and variances, and an MCMC algorithm to sample from the posterior distribution. The model is flexible enough to allow for both non-normality in the error distribution and serial correlation in the error variances.¹⁴

B.1 Description of the Model

One can rewrite the VAR model as a multivariate regression

$$y_t = X_t B + \epsilon_t, \quad (7)$$

$1 \times n$ $1 \times k$

where X_t includes the lagged values of y_t and a constant term. The error term ϵ_t has multivariate normal distribution with mean μ_{S_t} and variance-covariance matrix Σ_{S_t} , where S_t is the state of the system at time t . Each period, the state takes one of m possible values. The state evolves over time according to the stationary first-order Markov finite state model. Thus, the probability of state j given that in the previous period the state was i is

$$P[S_t = j \mid S_{t-1} = i, S_u(u < t - 1), A] = \quad (8)$$

$$P[S_t = j \mid S_{t-1} = i, A] = p_{ij}, \quad (9)$$

where A denotes the model.

Define the transition matrix as $\mathbf{P} = [p_{ij}]$ and the series of the states \mathbf{S} . In the stationary first-order Markov finite state model, the unconditional distribution of the state is always the

¹⁴Since the seminal contribution of Hamilton (1989), the number of papers using Markov regime-switching in dynamic models has been enormous. The model described in this section is similar to that in Sims and Zha (2006).

stationary distribution π with the property $\pi\mathbf{P} = \pi$. Thus, one does not need to impose a prior density on the initial value of the state. Denote the mapping from \mathbf{P} to the stationary distribution $\pi(\mathbf{P}) = [\pi_1(\mathbf{P}), \dots, \pi_m(\mathbf{P})]$. The likelihood with the states fixed is

$$p(\mathbf{S} \mid \mathbf{P}, A) = \pi_{S_1}(\mathbf{P}) \prod_{i,j} p_{ij}^{n_{ij}} \quad (10)$$

where n_{ij} is the number of dates $t = 2 \dots T$ such that $S_{t-1} = i$ and $S_t = j$.

Conditional on the state being equal to j , and the lagged values of y_t , ϵ_t is a multivariate normal random variable with mean μ_j and variance-covariance matrix Σ_j (these are grouped in μ and Σ , to ease notation).

The observables density conditional on the state is given by

$$y_t \mid \{X_t, B, \mu, \Sigma, S_t = j\} \sim N(X_t B + \mu_j, \Sigma_j). \quad (11)$$

Note that the X_t can include a constant term or not. If it does, it is possible to identify the constant and the means of the errors by introducing a tighter prior centered around zero for the error means μ_j . The vectors of interest are the impulse response functions, which depend only on $B(L)$ and Γ_0^{-1} , and the calculations do not involve the constant term or the error means separately. Thus, even though it is possible to include and identify a constant term separately from the error means –and this is the approach taken in this project– it is not necessary to do so, and the convergence may be slightly faster without including it.

The joint distribution of the y_t (denote all the observations by \mathbf{y} and \mathbf{X} , and group the

probabilities of the states in π and the errors in ϵ) is thus given by

$$p[\mathbf{y} \mid \mathbf{X}, B, \mu, \boldsymbol{\Sigma}, \mathbf{S}, \pi] \quad (12)$$

$$\propto \prod_{t=1}^T |\Sigma_{S_t}|^{-\frac{1}{2}} \exp \left\{ -\frac{1}{2} (y_t - X_t B - \mu_{S_t}) \Sigma_{S_t}^{-1} (y_t - X_t B - \mu_{S_t})' \right\} \quad (13)$$

$$= \prod_{j=1}^m |\Sigma_j|^{-T_j/2} \exp \left\{ -\frac{1}{2} \sum_{t:S_t=j} (y_t - X_t B - \mu_j) \Sigma_j^{-1} (y_t - X_t B - \mu_j)' \right\} \quad (14)$$

$$= \prod_{j=1}^m |\Sigma_j|^{-T_j/2} \exp \left\{ -\frac{1}{2} \text{tr}(\Sigma_j^{-1} \hat{\Sigma}_j) \right\}, \quad (15)$$

where $\hat{\Sigma}_j = (y - \mathbf{X}B)'(y - \mathbf{X}B)|_{\text{for } t:S_t=j}$.

B.1.1 Rewriting the Model in Stacked Form

It is useful to rewrite the system in stacked form. First, define

$$\mathbf{Z}_{T \times m} = [z_1, \dots, z_T]' = [\delta(S_t, j)], \quad \mathbf{W}_{T \times (m+k)} = [\mathbf{Z} \ \mathbf{X}], \quad \Phi_{(m+k) \times n} = [\mu \ B].$$

Define:

$$\gamma = \text{vec}(\Phi) \quad (16)$$

$$\tilde{\mathbf{y}} = \text{vec}(\mathbf{y}) \quad (17)$$

$$\tilde{\mathbf{W}} = I \otimes \mathbf{W} \quad (18)$$

$$\tilde{\epsilon} = \text{vec}(\epsilon) \quad (19)$$

$$\mathbf{Q} = \text{Cov}(\tilde{\epsilon}) = \left[\sum_{j=1}^m (\Sigma_j \otimes \text{diag}(\mathbf{Z}_{\cdot,j})) \right]^{-1}. \quad (20)$$

With this notation, we can rewrite the joint distribution of \tilde{y} as

$$p(\tilde{y} \mid \mathbf{X}, \gamma, \boldsymbol{\Sigma}, \mathbf{S}, \pi) \propto |\mathbf{Q}|^{1/2} \exp \left[-\frac{1}{2}(\tilde{y} - \tilde{W}\gamma)' \mathbf{Q}(\tilde{y} - \tilde{W}\gamma) \right].$$

B.1.2 Prior

The prior for each column \mathbf{p}_i of the transition matrix \mathbf{P} is Dirichlet with parameters a_{ij} :

$$p(\mathbf{p}_i) \propto \prod_{j=1}^m p_{ij}^{a_{ij}-1}. \quad (21)$$

For the Σ_j , it is convenient to use independent inverse Wishart priors:

$$p(\Sigma_j^{-1}) \propto |\Sigma_j|^{-\frac{1}{2}(\nu-2)} \exp \left(-\frac{1}{2} \text{tr} \Sigma_j^{-1} \underline{H} \right), \quad (22)$$

that is, Σ_j^{-1} is assumed to have a Wishart distribution with $\nu + n - 1$ degrees of freedom with scale matrix \underline{H} .

Finally, the prior for γ is a multivariate normal with mean $\underline{\gamma}$ and precision matrix \underline{H}_γ , where:

$$\underline{\gamma}_{n(m+k) \times 1} = \begin{pmatrix} \underline{0} \\ \underline{\beta} \end{pmatrix}, \quad \underline{H}_\gamma_{n(m+k) \times n(m+k)} = \begin{bmatrix} h_\alpha \mathbf{I}_{nm} & 0 \\ 0 & \underline{H}_\beta \end{bmatrix}.$$

There are two ways to implement a Minnesota prior in this framework. The first is to incorporate it directly into the prior for γ . The second approach –the one used in this paper– is to add the dummy observations to the dataset as described in Sims and Zha (1998). Then, one can simply work with the markov-switching VAR model just described using the dummy-observation-augmented data. The parameters for the dummies are the same as those described in Table A.1.

In addition to the dummy observations, I set values for the prior parameters in the model.

The Dirichlet parameters of the prior for \mathbf{P} are all set to 1. For the prior on Σ_j , I use ν equal to 10, which corresponds to $\nu + n - 1 = 11$ degrees of freedom and scale matrix equal to the inverse of the variance covariance matrix of the errors in the results from a BVAR using Sims' codes with the same data, divided by $\nu + n - 1$. The coefficients are (after thinking of the dummies as data) all centered around zero. The precision matrix is diagonal, with a precision of zero for the constant and the coefficients on the lagged variables, and a precision of 1 for the error means, which allows identification. Despite some elements of the diagonal of the precision matrix being set to zero, the prior is not improper because of the inclusion of the dummy observations. The number of mixtures is fixed at five and, as before, I use 16 lags.

B.2 MCMC Algorithm for Exploring the Posterior Distribution

The Metropolis-within-Gibbs algorithm to sample from the posterior is based on making successive draws from the conditional posterior distributions of γ , Σ_j , \mathbf{P} , and the time series of states S_t . In the stationary first-order Markov finite state model there is a restriction on \mathbf{P} and thus the posterior is not a standard distribution. However, sampling is possible using an independence Metropolis-Hastings step.

The conditional posterior density of γ is given by

$$\begin{aligned} \gamma & \mid (\mathbf{X}, \tilde{y}, \Sigma, \mathbf{S}, \pi) \sim N(\bar{\gamma}, \bar{H}_\gamma) \\ \bar{H}_\gamma & = \underline{H}_\gamma + \tilde{W}\mathbf{Q}\tilde{W} \\ \bar{\gamma} & = \bar{H}_\gamma^{-1} \left[\underline{H}_\gamma \underline{\gamma} + \tilde{W}\mathbf{Q}\tilde{y} \right]. \end{aligned}$$

This is analogous to the derivation in Geweke (2005), section 6.4.2. Note that, because \mathbf{Q} is not diagonal, the system cannot be estimated equation by equation as in the standard Vector Autoregression model.

The conditional posterior density of Σ_j is given by the product of (13) and (22), which is an inverse Wishart with $T_j + \nu + n - 1$ degrees of freedom and scale matrix $[\hat{\Sigma}_j + \underline{H}^{-1}]^{-1}$:

$$\begin{aligned} \Sigma_j &| (\mathbf{X}, \tilde{\mathbf{y}}, \gamma, \mathbf{S}, \pi) \\ &\propto |\Sigma_j|^{-T_j/2} \exp\left\{-\frac{1}{2}\text{tr}(\Sigma_j^{-1}\hat{\Sigma}_j)\right\} |\Sigma_j|^{-\frac{1}{2}(\nu-2)} \exp\left(-\frac{1}{2}\text{tr}\Sigma_j^{-1}\underline{H}^{-1}\right) \\ &= |\Sigma_j|^{-(T_j+\nu-2)/2} \exp\left\{-\frac{1}{2}\text{tr}(\Sigma_j^{-1}(\hat{\Sigma}_j + \underline{H}^{-1}))\right\}. \end{aligned}$$

The conditional posterior of the states is

$$\begin{aligned} p(\mathbf{S} | \gamma, \Sigma, \mathbf{P}, \mathbf{y}, \mathbf{X}, A) &\propto \pi_{S_1} |\Sigma_{S_1}|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}(y_1 - X_1 B - \mu_{S_1}) \Sigma_{S_1}^{-1} (y_1 - X_1 B - \mu_{S_1})'\right] \\ &\prod_{t=2}^T p_{S_{t-1}S_t} |\Sigma_{S_t}|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}(y_t - X_t B - \mu_{S_t}) \Sigma_{S_t}^{-1} (y_t - X_t B - \mu_{S_t})'\right]. \end{aligned}$$

An algorithm to sample directly from this density is due to Chib (1996).

The conditional posterior density of \mathbf{P} is

$$p(\mathbf{P} | \mathbf{S}, A) = \pi_{S_1}(\mathbf{P}) \prod_{i,j} p_{ij}^{a_{ij}+n_{ij}-1}. \quad (23)$$

As described in Geweke (2005), section 7.2.2, a method to sample from this distribution is using an independence Metropolis-Hastings algorithm. First, each column \mathbf{p}_i^* of the candidate \mathbf{P}^* is drawn from a Dirichlet with parameters $a_{i1} + n_{i1} - 1, \dots, a_{im} + n_{im} - 1$. The acceptance probability is $\min[\pi_j(\mathbf{P}^*)/\pi_j(\mathbf{P}), 1]$.

B.3 Results

Impulse responses from the Markov-Switching model are shown in Figure B.1. The responses look similar to those from the standard BVAR, although the estimated effect of interest in

electric cars on oil prices is smaller. Figure B.2 shows the estimated variances over time from the Markov-Switching model, calculated as the average variance for each time period over all the iterations of the posterior sampling. Interest in electric cars has several periods of high variance; the variance of oil prices, on the other hand, is quite stable except for a period of high variance from September 2008 until January 2009. Figure B.3 shows the estimated density of the stationary distribution in the last iteration of the Gibbs sampler, together with the estimated residuals from that iteration.

C Interest in Other Alternatives

In this section, I compare the response of oil prices to interest in electric cars with the response to interest in hybrid and alternative vehicles in general, interest in hybrids in particular, and interest in ethanol. To measure interest in hybrid and alternative vehicles in general, I use the “Hybrids and Alternative Vehicles” category of Google Insights. This category captures a broad range of search terms related to alternative vehicles, including electric cars and hybrids. Using Google Insights, I find that all the top ten search terms related to “hybrids” in the automotive category include the words “hybrid” or “hybrids”. All the top ten searches related to “ethanol” include the terms “ethanol” or “e85”.

Figure C.1 shows the responses of oil prices to the search volume indexes for electric cars, hybrid and alternative vehicles, hybrids, and ethanol. The responses to the hybrid and alternative vehicles category and to the search index for hybrids are similar to the response to interest in electric cars, although the latter response is stronger. The three measures of search volume are highly correlated (the correlation between any of them is around 90%). The response to interest in ethanol is also negative, but the error bands are much wider. The correlation between interest in ethanol and the other three series is lower, around 50%.

D Estimation Without the Recession Period

Figure D.1 shows the impulse responses for the pre-recession period. This excludes from the sample 86 weeks of data, from December 2007 to July 2009. The magnitude of a typical oil shock was considerably smaller for the pre-recession period, as was the response to a shock to interest in electric cars. The response to interest in electric cars is still of a magnitude comparable to a typical oil shock.

References

- Aruoba, S. Boragan, Francis X. Diebold, and Chiara Scotti**, “Real-Time Measurement of Business Conditions,” Working Paper 08-19, Federal Reserve Bank of Philadelphia May 2008.
- Barberis, Nicholas, Andrei Schleifer, and Robert Vishny**, “A Model of Investor Sentiment,” *Journal of Financial Economics*, September 1998, *49*, 307–343.
- Becker, Thomas A.**, “Electric Vehicles in the United States: A New Model With Forecasts to 2030,” Center for Entrepreneurship and Technology Technical Brief 2009.1.v1.v1, University of California, Berkeley July 2009.
- Better Place**, “An Economic and Environmental Blueprint for the Future of Energy and Transportation,” Technical Report, Better Place 2008.
- British Petroleum**, *BP Statistical Review of World Energy June 2009*, British Petroleum, June 2009.
- Chib, Siddharta**, “Calculating Posterior Distributions and Modal Estimates in Markov Mixture Models,” *Journal of Econometrics*, 1996, *95*, 79–97.

- Choi, Hyunyoung and Hal Varian**, “Predicting the Present with Google Trends,” Technical Report, Google Inc. April 2009.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers**, “Speculative Dynamics and the Role of Feedback Traders,” *The American Economic Review*, 1990, *80* (2), 63–68.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam**, “Investor Psychology and Security Market under- and Overreactions,” *The Journal of Finance*, 1998, *53* (6), 1839–1885.
- Deutsche Bank**, “Electric Cars: Plugged In,” Technical Report, Deutsche Bank, Global Markets Research 2008.
- Dvir, Eyal and Ken Rogoff**, “The Three Epochs of Oil,” Boston College Working Papers in Economics 706, Boston College Department of Economics April 2009.
- Geweke, John**, *Contemporary Bayesian Econometrics and Statistics*, Hoboken, NJ: John Wiley and Sons., 2005.
- Ginsberg, Jeremy, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant**, “Detecting influenza epidemics using search engine query data,” *Nature*, 2009, *457*, 1012–1014.
- Hamilton, James D.**, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 1989, *57* (2), 357–384.
- , “Causes and Consequences of the Oil Shock of 2007-08,” Working Paper 15002, National Bureau of Economic Research May 2009.
- , “Understanding Crude Oil Prices,” *Energy Journal*, 2009, *30* (2), 179–206.

Hong, Harrison and Jeremy C. Stein, “A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets,” *Journal of Finance*, December 1999, *54* (6), 2143–2184.

– , **Terence Lim, and Jeremy C. Stein**, “Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies,” *The Journal of Finance*, 2000, *55* (1), 265–295.

Jegadeesh, Narasimhan and Sheridan Titman, “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *The Journal of Finance*, 1993, *48* (1), 65–91.

Kilian, Lutz, “The Economic Effects of Energy Price Shocks,” *Journal of Economic Literature*, 2008, *46* (4), 871–909.

Polgreen, Philip M., Yiling Chen, David M. Pennock, and Forrest D. Nelson, “Using Internet Searches for Influenza Surveillance,” *Clinical Infectious Diseases*, 2008, *47* (11), 1443–1448.

Sims, Christopher A. and Tao Zha, “Bayesian Methods for Dynamic Multivariate Models,” *International Economic Review*, November 1998, *39* (4), 949–68.

– **and** – , “Error Bands for Impulse Responses,” *Econometrica*, September 1999, *67* (5), 1113–1156.

– **and** – , “Were There Regime Switches in U.S. Monetary Policy?,” *The American Economic Review*, 2006, *96* (1), 54–81.

Williamson, Harold F. and Arnold R. Daum, *The American Petroleum Industry: The Age of Illumination 1859-1899.*, New York, NY: Northwestern University Press, 1959.

	Electricity	Gasoline			
		US	Brazil	UK	Germany
Cost per Gallon / kWh	0.10	4.00	5.95	8.38	9.28
Miles per Gallon / kWh	5	35	35	35	35
Fuel Cost per Mile	0.02	0.11	0.17	0.24	0.27
Battery Depr per Mile	0.08	-	-	-	-
Miles per Year	15,000	15,000	15,000	15,000	15,000
Fuel Cost per Year	1,500	1,714	2,550	3,591	3,977

Table 1
Cost of Fueling: Electric vs. Gasoline (\$)

Source: Deutsche Bank (2008).

Parameter	Description	Value
λ	Weight on the sum-of-coefficients dummy	5
μ	Weight on the no-cointegration dummies	2
π_1	Tightness parameter on the ‘Minnesota prior dummies’	3
π_2	Weight of own lags relative to coefficients on other variables	1
π_3	Decay parameter	1/2
σ_k	Degree of variability of the k th variable	$stdev(\Delta y_k)$
w	Weight on variance-covariance dummy	1

Table A.1
Parameter Values for the Minnesota Prior

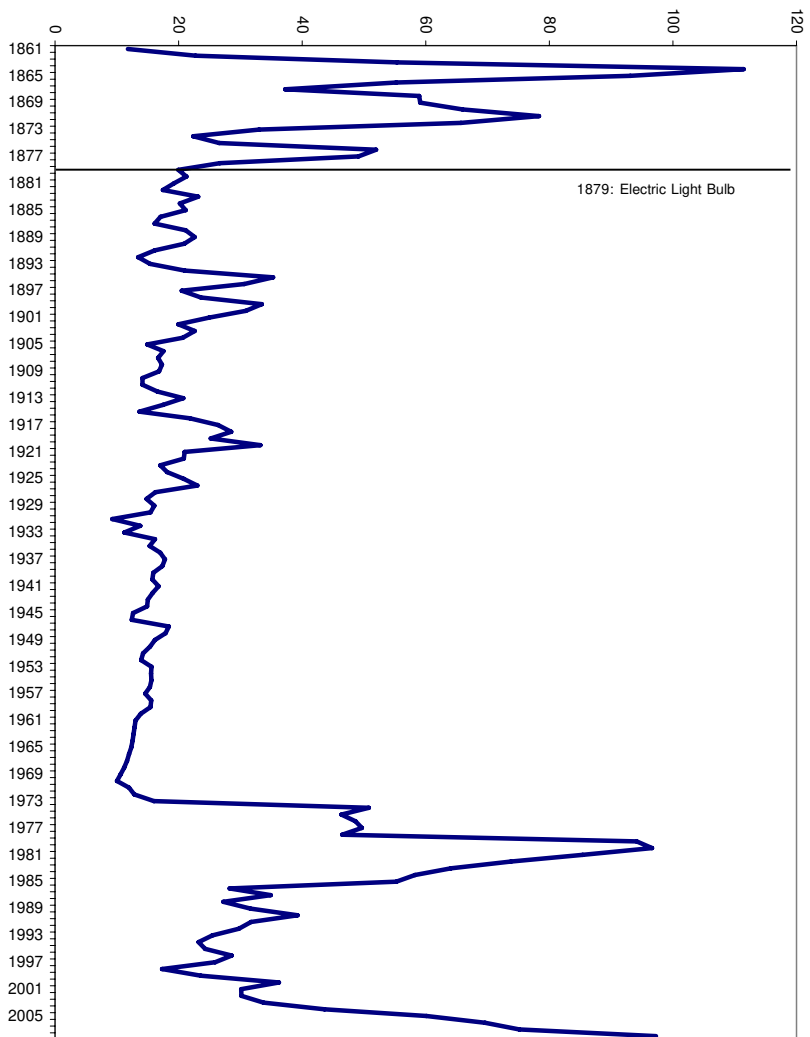


Figure 1
Real Oil Prices Since 1861 in Constant 2008 US Dollars

Source: British Petroleum (2009). Vertical line added by the author.

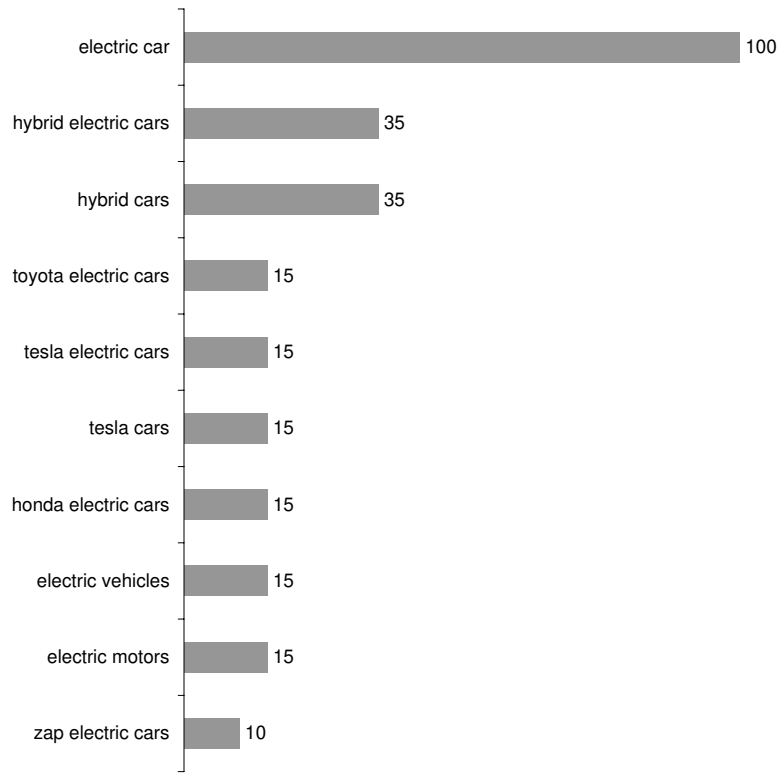




Figure 2
 Top Ten Search Terms Related to Electric Cars, 2004-2009


Source: Google Insights. This figure shows the keywords that users search for the most before searching for “electric cars,” as well as after. The numbers measure the relative intensity of searches for the phrases before and after searches for “electric cars.” The volume numbers are normalized, with the top search assigned a volume of 100.



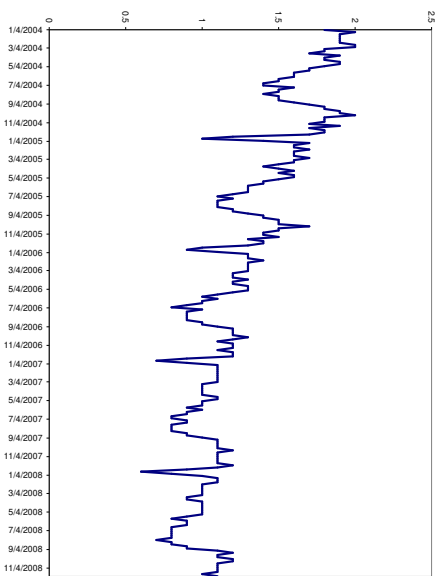
(a) Ice Cream



(b) Pumpkin



(c) Down Goose



(d) Economics

Figure 3
 Google Search Volume Index for "Ice Cream", "Pumpkin", "Down Goose", and
 "Economics"




Figure 4
Interest in Electric Cars and Oil Prices

Source: Google Trends, Energy Information Administration, and author's calculations.

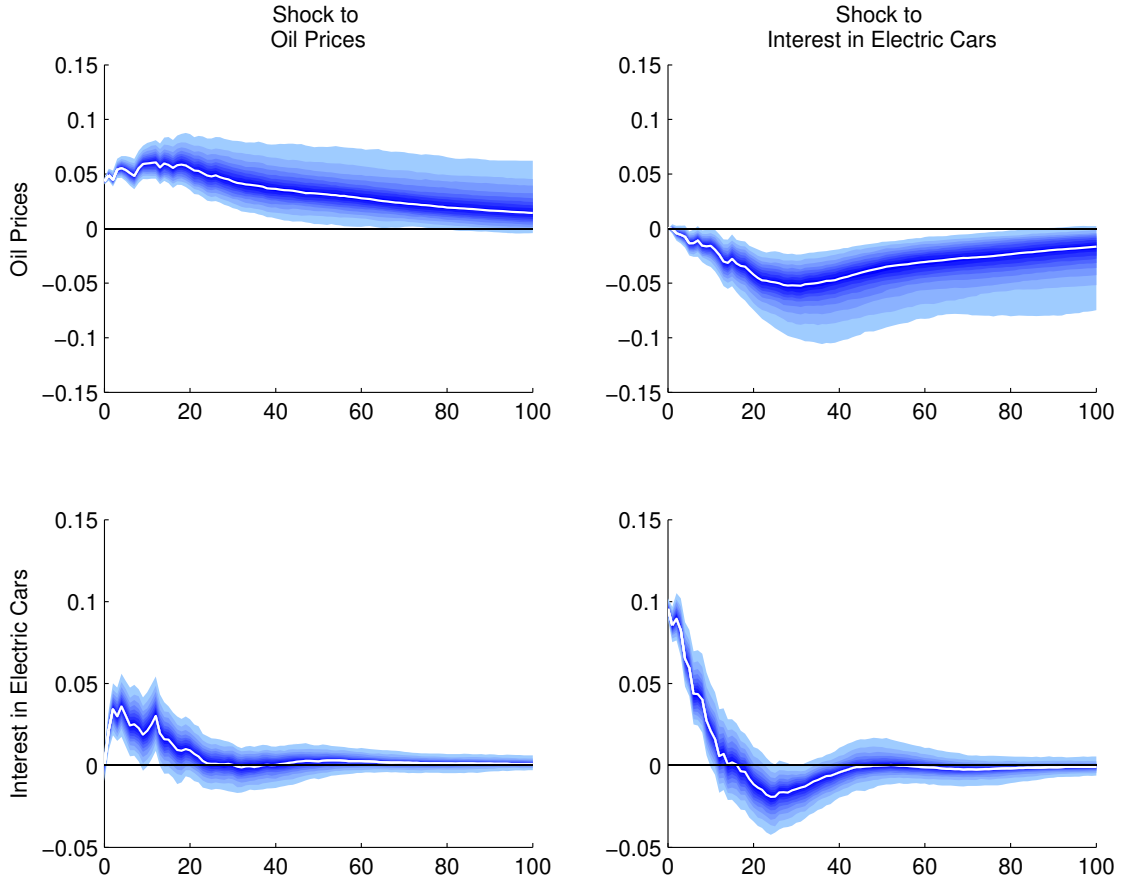


Figure 5
 Posterior Density of Impulse Responses to Shocks to Interest in Electric Cars and Oil Prices

Note: In every case, the x axis indicates the number of weeks after the shock. The first column shows the response of each variable to a shock to oil prices and the second the response to a shock to interest in electric cars.




Figure 6
Simulation Without Shocks






Figure 7
Simulation Using Only Shocks to Interest in Electric Cars, Subtracting the Values
Predicted Without Shocks



(a) Shocks to interest in electric cars.



(b) Shocks to oil prices.

Figure 8
Quantile-Quantile Plots of the Residuals

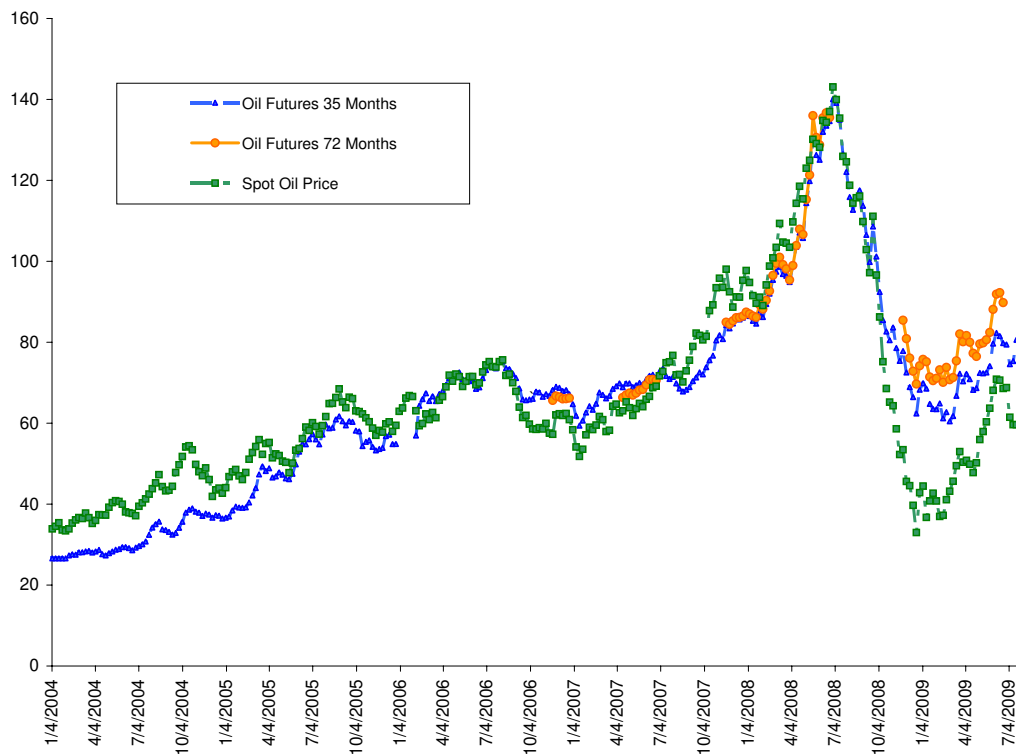


Figure 9
Oil Spot and Futures Prices

Note: Weekly average of closing prices in US dollars. *Source:* Bloomberg, Energy Information Administration, and author's calculations.

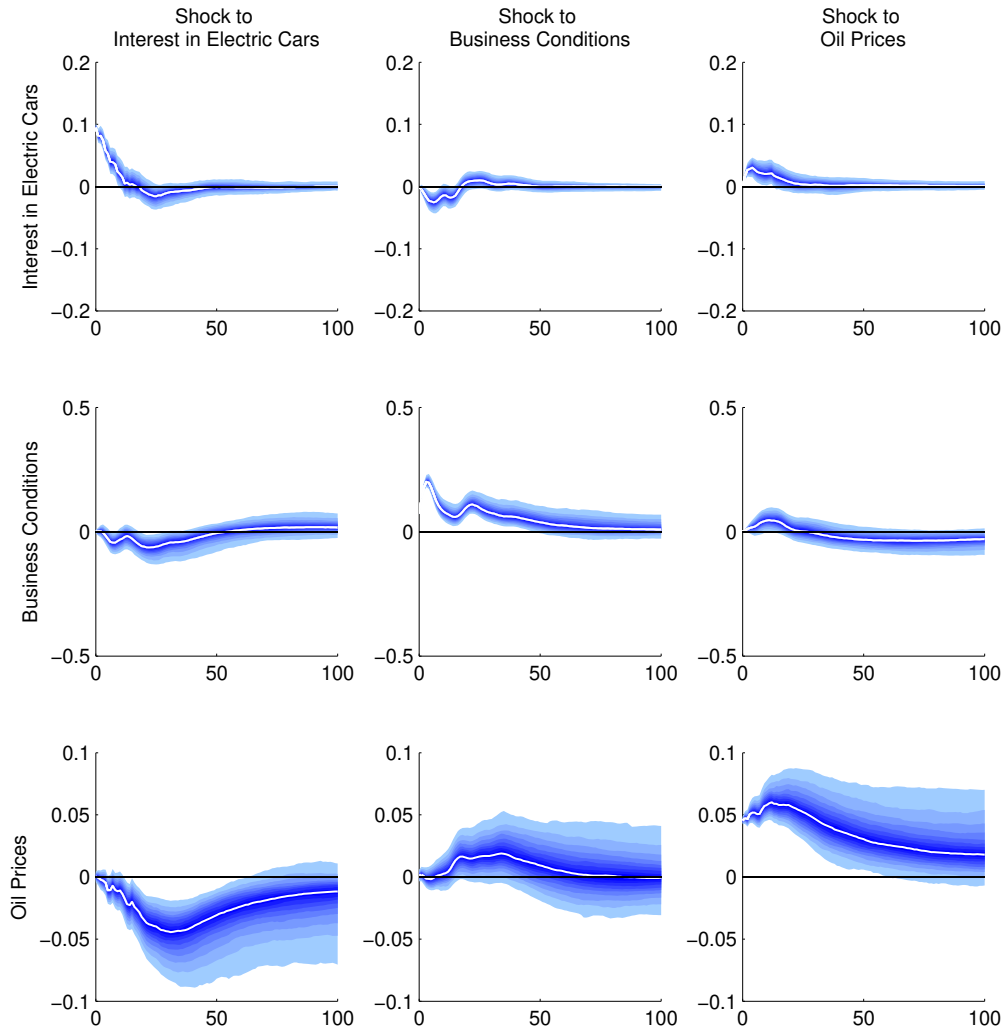

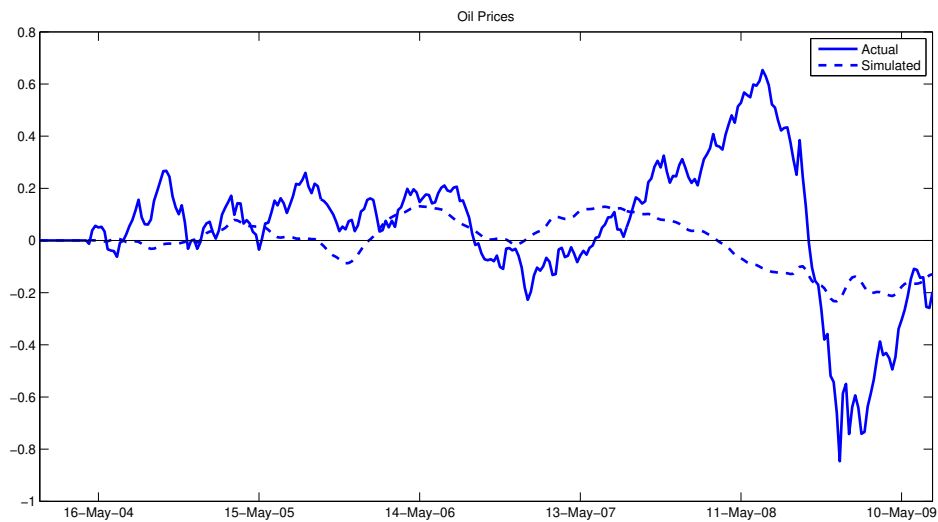


Figure 10
 Three-Variable BVAR: Posterior Density of the Impulse Responses

Note: In every case, the x axis indicates the number of weeks after the shock.




(a) Only Shocks to Interest in Electric Cars




(b) Only Shocks to Business Conditions


Figure 11
 Three-Variable BVAR: Simulations Using Only Shocks to Interest in Electric Cars, and
 Using Only Shocks to Business Conditions, Subtracting the Values Predicted Without
 Shocks



(a) $\lambda = 5$




(b) $\lambda = 0$




(c) $\lambda = 10$


Figure A.1
 Responses for Different Values of λ , Holding Other Hyperparameters Fixed at the Baseline Values



(a) $\mu = 2$




(b) $\mu = 0$




(c) $\mu = 4$


Figure A.2
Responses for Different Values of μ , Holding Other Hyperparameters Fixed at the Baseline Values



(a) $\pi_1 = 3$

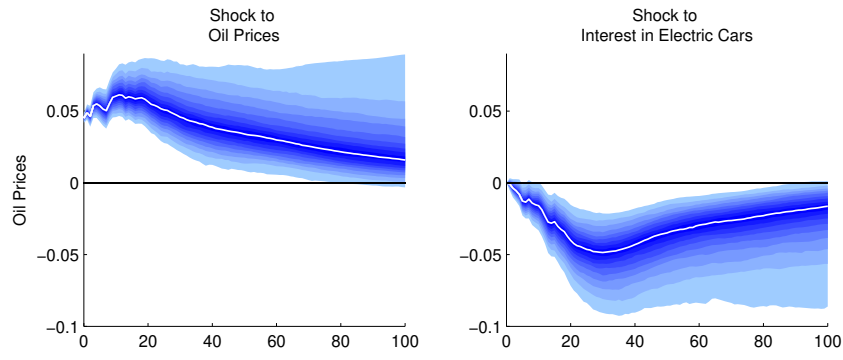


(b) $\pi_1 = 0$




(c) $\pi_1 = 6$


Figure A.3
Responses for Different Values of π_1 , Holding Other Hyperparameters Fixed at the Baseline Values



(a) $\pi_3 = 1/2$



(b) $\pi_3 = 0$



(c) $\pi_3 = 1$

Figure A.4
Responses for Different Values of π_3 , Holding Other Hyperparameters Fixed at the Baseline Values

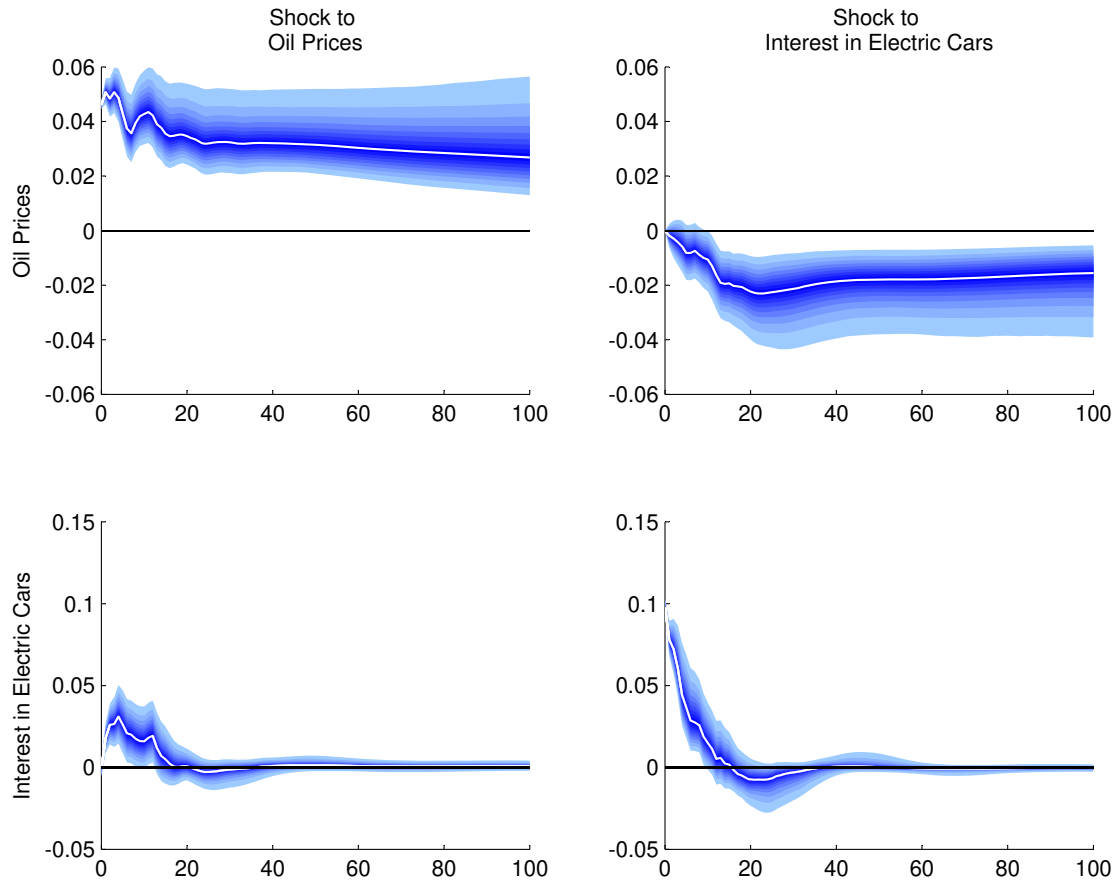


Figure B.1
Impulse Responses from the Markov Switching BVAR




Figure B.2
 Estimated Variances of the Shocks Over Time, and Estimated Residuals Squared

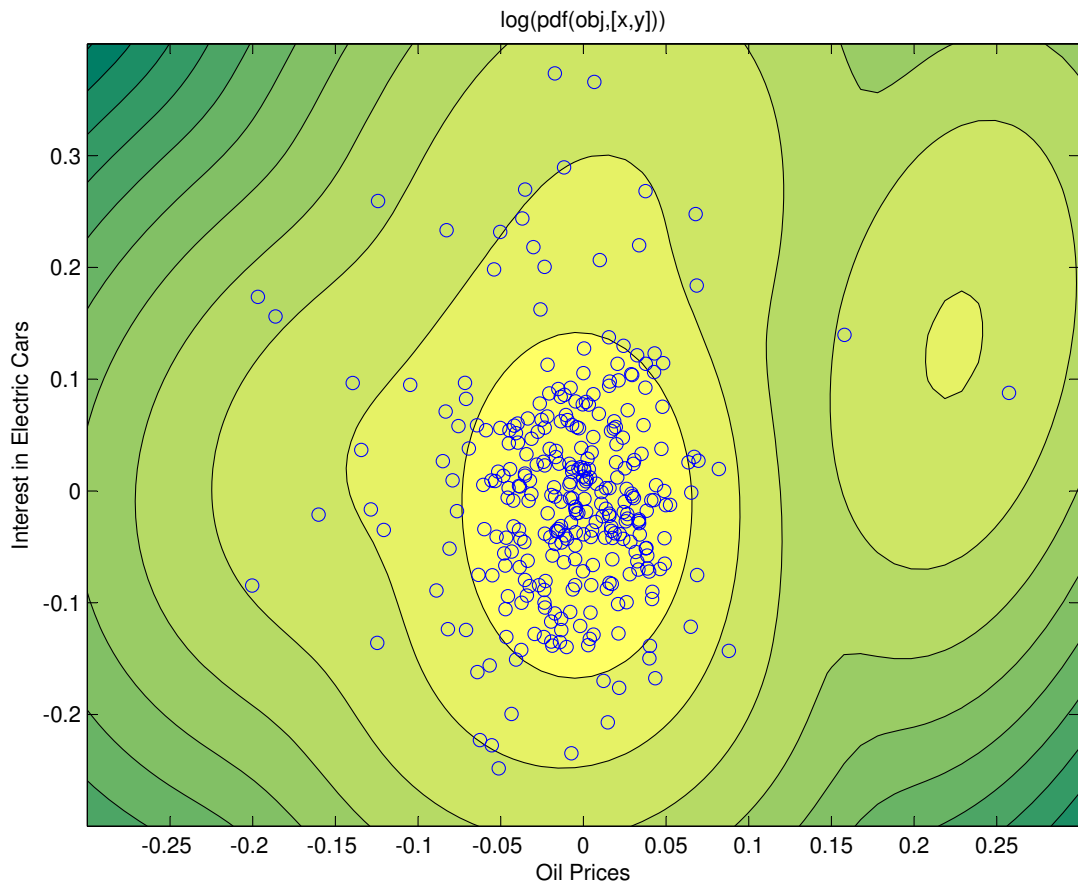
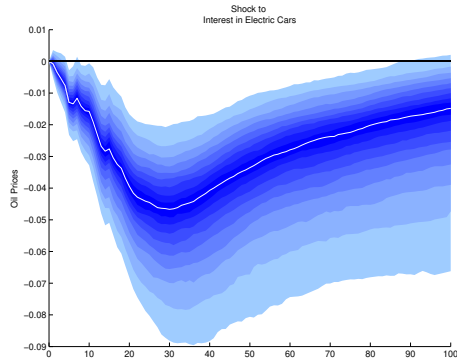




Figure B.3
Estimated Shocks and Contour plot of the Estimated Joint Density of the Shocks (Last Iteration of the MCMC Simulation)




(a) Electric Cars



(b) 'Hybrid and Alternative Vehicles'



(c) Hybrids



(d) Ethanol

Figure C.1
Responses of Oil prices to a Shock to Interest in Different Alternatives

Note: Two-variable BVARs estimated separately, including prices and the given search phrase search volume index.

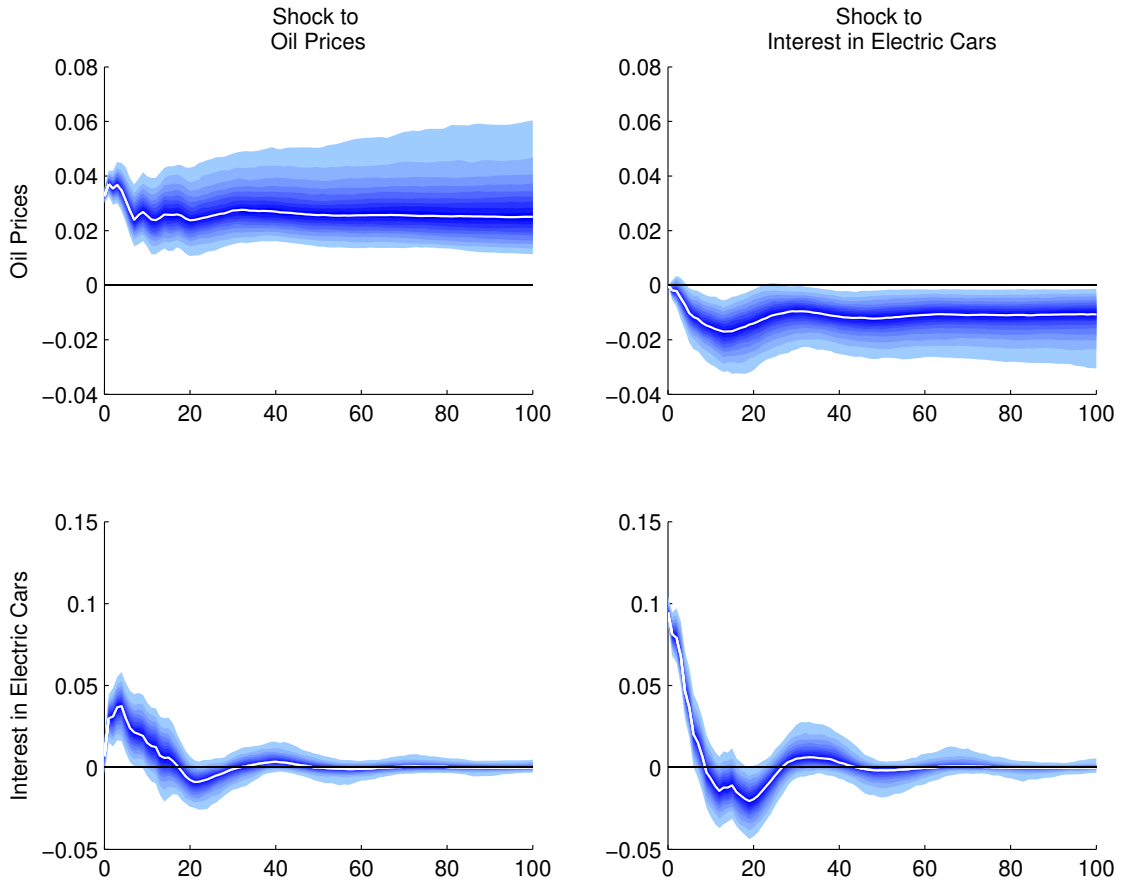


Figure D.1
 Impulse Responses Without the Recession Period (January 2004 - November 2007)