Jo, Soojin. 2014. "The Effects of Oil Price Uncertainty on Global Real Economic Activity." *Journal of Money, Credit and Banking* 46, no. 6: 1113–35.

Kilian, Lutz, and Robert J. Vigfusson. 2011. "Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?" *Quantitative Economics* 2, no. 3: 419–53.

Kim, Dong Heon. 2012. "What Is an Oil Shock? Panel Data Evidence." *Empirical Economics* 43, no. 1: 121–43.

Lee, Kiseok, Shawn Ni, and Ronald A. Ratti. 1995. "Oil Shocks and the Macroeconomy: The Role of Price Variability." *Energy Journal* 16, no. 4: 39–56.

Loungani, Prakash. 1986. "Oil Price Shocks and the Dispersion Hypothesis." *Review of Economics and Statistics* 68, no. 3: 536–39.

Mork, Knut Anton. 1989. "Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results." *Journal of Political Economy* 97, no. 3: 740–44.

Ravazzolo, Francesco, and Philip Rothman. 2013. "Oil and U.S. GDP: A Real-Time Out-of-Sample Examination." *Journal of Money, Credit and Banking* 45, nos. 2–3: 449–63.

U.S. Energy Information Administration. 2016. *Monthly Energy Review,* October. Washington: U.S. Department of Energy.

COMMENT BY

VALERIE A. RAMEY This paper by Christiane Baumeister and Lutz Kilian represents a very interesting and informative analysis of the likely effects of the recent oil price decline on the U.S. economy. It is widely believed that past instances of dramatic oil price rises were detrimental to the U.S. economy. A question that arose when oil prices plummeted in 1986—and that has arisen again with the more recent collapse of oil prices—is how much do *falling* oil prices *help* the U.S. economy. This question has become even more interesting in the face of the recent oil shale revolution, because there are now more industries that may lose when oil prices fall.

Baumeister and Kilian present a comprehensive analysis of the many potential channels of transmission that have been discussed in the oil shock literature. For example, they consider (i) the cost channel, through which changes in oil prices affect production costs; (ii) the sectoral reallocation channel, which lowers GDP due to reallocation frictions; (iii) the uncertainty channel, through which changes in uncertainty about future oil prices affect current spending; (iv) the operating cost channel, through which the price of oil affects the demand for goods that are complements in consumption, such as motor vehicles; and (v) the discretionary income

channel, through which a fall in oil prices gives consumers more discretionary income, which they then spend on other goods. On the basis of many data series that they collect and inspect, Baumeister and Kilian argue that there is not much evidence that most of these channels have a significant effect. Their main finding is that the discretionary income effect—operating on both consumption and non-oil, nonresidential investment—boosted GDP, but this effect was mostly counteracted by a dramatic decline in oil sector investment spending.

With the exception of the analysis of the discretionary income effect, Baumeister and Kilian rely almost exclusively on informative case study analyses that compare effects across time (such as the period from 1986 to the current episode), across states (oil-producing versus non-oil-producing), and across categories of expenditures. There is much to be learned from the detailed data and patterns that the authors present in their case studies. It is important to note, however, that all their assessments depend on their nonstatistical judgment about what the counterfactual should have been. For example, when assessing the importance of frictions in the reallocation of labor in their subsection III.A, the authors argue, "If frictions in reallocating labor drove up unemployment after June 2014, this would imply that—in the absence of these frictions—unemployment would have dropped even more sharply than it actually did, which does not seem plausible." Thus, their counterfactual is based on their assessment of plausibility. My inspection of the time series reveals that the unemployment rate fell faster during several other periods, such as the early 1980s and mid-1990s. Thus, it is not clear to me whether a faster decline is implausible.

I focus the rest of my discussion on two main points. First, I argue that the discretionary income hypothesis, which has been discussed in several places in the literature and on which Baumeister and Kilian base their consumption calculations, makes no economic sense. I argue that the literature has too often confused the terms-of-trade effect, which has a sound economic basis, with this discretionary income effect, which has no economic foundation. Second, I present arguments and evidence for why we should expect a given oil price change to have a smaller effect now than in earlier decades.

THE DISCRETIONARY INCOME EFFECT VERSUS THE TERMS-OF-TRADE EFFECT

I begin by discussing the discretionary income effect, also referred to as the consumer purchasing power effect (Edelstein and Kilian 2009). According to Baumeister and Kilian, "In these models, a drop in the real retail price of gasoline is akin to a tax cut from the point of view of consumers, which is expected to stimulate private consumption and hence real GDP." They refer

to papers in the literature, such as the one by Paul Edelstein and Kilian (2009), and a quotation from a speech by Janet Yellen (2011). Edelstein and Kilian (2009, pp. 766–67), who present no theoretical model, state in the text, "Higher energy prices are expected to reduce discretionary income, as consumers have less money to spend after paying their energy bills." Edelstein and Kilian (2009) mention in a footnote that this effect must work through net imports and cannot work if all oil is domestically produced; yet in the rest of the paper, they proceed as though the effect has nothing to do with net imports. In particular, they scale gasoline price changes by the gasoline share of consumption expenditures and do not mention the import share of petroleum. In contrast, Yellen's (2011) speech, which is partly quoted by the authors in their introduction, focuses entirely on the terms-of-trade effect, which has good economic foundations. However, it is unrelated to Edelstein and Kilian's (2009) notion or to Baumeister and Kilian's discussion (see the beginning of their subsection II.B).

Perhaps the best way to illustrate why the discretionary income channel makes no sense is in the context of Baumeister and Kilian's back-of-the-envelope calculation. They assess the impact of the 45 percent decline in the real price of gasoline over the last two years. They use an estimate of a –0.37 short-run elasticity of gasoline demand and the fact that gasoline expenditures are 3.17 percent of total consumption expenditures to argue that the decline in the real price of gasoline freed up 1.13 percent of income for additional consumption purchases. However, the same logic would imply that the decline in the relative price of any good that had a demand elasticity less than −1 would act like a tax cut and stimulate consumption and GDP. For example, consider the behavior of the services component of consumption expenditures in 2015. The real price of consumer services rose 1.7 percent. This is a much smaller price change than gasoline prices, but services consumption accounts for 67 percent of consumption expenditures. If I use Baumeister and Kilian's calculation (and assume a very low demand elasticity for this large category), I find that this change reduced discretionary income by 1.1 percent, which by coincidence is identical to the amount by which gasoline prices supposedly increased it. The point of my example is to show that relative price changes should not have direct effects on aggregate consumption, independent of their effect on income, because a decrease in the relative price of one good means a corresponding increase in the relative price of another good.

A relative price change can, however, have effects on consumption in an open economy. This is the terms-of-trade mechanism to which Yellen

(2011) refers, and it is the effect that is used in the calculations in the 2016 *Economic Report of the President.* As the discussion in box 2-1 of the report points out, the effect of a decline in the price of a good for which the United States is a net importer works directly through a boost in real domestic income. David Backus and Mario Crucini (2000) show that oil price fluctuations accounted for a significant part of the terms-of-trade fluctuations of the United States during the 1970s and early 1980s.

To see this effect in the simplest possible model, consider a small, open economy that produces and consumes two goods, oil and non-oil. Because the country is open to trade, it is not required to consume the same amount that it produces of each good. Assuming that the current account is zero, the country faces the following budget constraint:

$$P_{oil}Y_{oil} + P_{non\text{-}oil}Y_{non\text{-}oil} = P_{oil}C_{oil} + P_{non\text{-}oil}C_{non\text{-}oil},$$

where P denotes the prices of each good, Y denotes the quantities produced of each good, and C denotes the quantities consumed of each good.

We can rewrite this budget constraint in a way that makes it clear how import status and relative prices interact:

$$C_{non\text{-}oil} = Y_{non\text{-}oil} + \frac{P_{oil}}{P_{non\text{-}oil}}(Y_{oil} - C_{oil}). \tag{1}$$

The term in parentheses is the difference between the country's production of oil and consumption of oil. If the country is a net exporter of oil ($Y_{oil} - C_{oil} > 0$), then a rise in the relative price of oil will raise non-oil consumption through the terms-of-trade effect on real income. If, however, the country is a net importer of oil ($Y_{oil} - C_{oil} < 0$), then a rise in the relative price of oil will lower non-oil consumption. The United States is a net importer of oil, so this simple model implies that non-oil consumption moves inversely with the relative price of oil. This terms-of-trade effect operates through an income effect.

This simple example also highlights two other important features. First, the correct scaling factor for looking at the effect of an oil price shock is not Edelstein and Kilian's (2009) consumption expenditure share on oil, but the net import status of the United States. Second, as the United States' net import status changes, we should expect the impact of a given change in oil prices to change as well. I return to these points below.

In response to the conference version of my discussion, Baumeister and Kilian have added a section in which they extend Edelstein and Kilian's

Table 1. Effect of an Oil Shock on Consumption[a]

Sample	*Estimate*
February 1970–December 1992	11.40
	(1.88)
January 1993–March 2016	2.58
	(1.20)

Source: My calculations, using Baumeister and Kilian's data.

a. My regression differs from Baumeister and Kilian's, and is described in the text. Heteroskedasticity- and autocorrelation-consistent standard errors are in parentheses.

(2009) scaling of shocks to include an additional multiplicative factor with the share of imports in total U.S. gasoline consumption. This is certainly a step in the right direction, but I would argue that Edelstein and Kilian's (2009) original term for the gasoline share of consumption expenditures should not even appear in this equation.

IS THE RELATIONSHIP BETWEEN OIL PRICES AND THE REAL ECONOMY STABLE?

The second issue concerns the stability of the relationship between oil price shocks and the economy. Olivier Blanchard and Jordi Galí (2010) present evidence that there was a structural break in the relationship between oil and the U.S. economy in the mid-1980s, with the impact of a given change in oil prices being much larger in the earlier period. Edelstein and Kilian (2009) show the same type of evidence when they use their purchasing power shocks (see their figure 4). Daniel Vine and I (2011) show that several other leading measures of oil shocks also suggest a difference in impact across periods.

To see if there might be a problem with the stability of Baumeister and Kilian's econometric equation that estimates the impact of their scaled price shock on consumption, I reestimated the equation, allowing the key coefficients to change halfway through the sample (in 1993). I make a few minor changes to the specification, such as including the log levels of real consumption rather than differences, and in order to summarize the cumulative effect on log consumption through 20 months (the integral under the impulse response function) with just one coefficient, I redefine the dependent variable as the integral of log real consumption from horizon 0 to horizon 20, and use Òscar Jordà's (2005) local projection method. My table 1 shows that the 1993–2016 sample coefficient is less than one-quarter of the 1970–92 sample coefficient. The p value for the hypothesis test that the coefficients are the same is 0.000. Thus, the estimates strongly suggest that oil price shocks have a much smaller effect on consumption now than in the 1970s and 1980s.

Figure 1. Terms of Trade and Crude Oil Prices, 1947–2016[a]

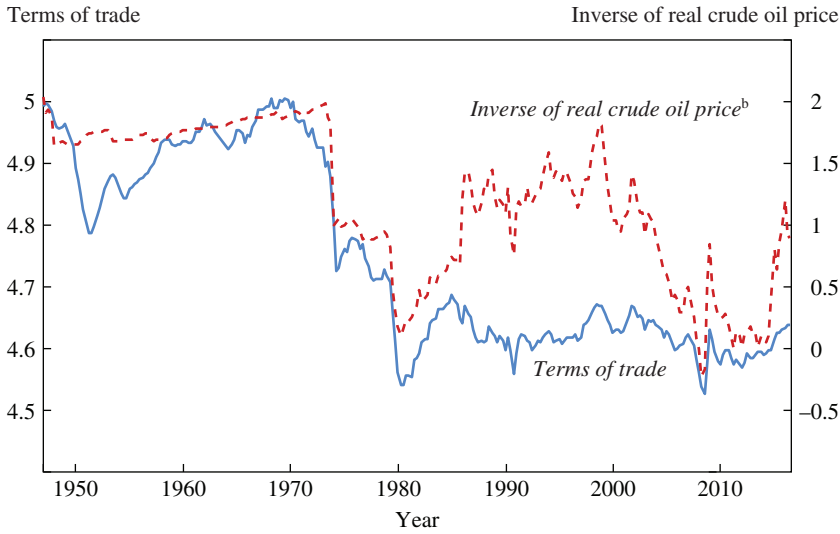

Source: Federal Reserve Economic Data.
a. Both series are on a log scale.
b. The crude oil price is the West Texas Intermediate spot price, reduced by the GDP deflator.

In fact, there are several reasons why we should expect oil price shocks to have much less of an effect on the real economy than they did in the 1970s. First, the terms-of-trade effects of oil shocks were much larger in the 1970s, as noted by Backus and Crucini (2000). My figure 1 shows the U.S. terms of trade and the inverse of real crude oil prices for the post–World War II period. The two big oil shocks of the 1970s led to a dramatic deterioration in the U.S. terms of trade. In contrast, the effects of recent decades' dramatic oil price changes on the terms of trade are barely evident in the figure.

Second, as argued by Ramey and Vine (2011), the oil shocks of the 1970s were accompanied by price controls and shortages, so the actual shocks were much larger than what is captured by looking at oil price changes alone. My figure 2 shows an updated version of Ramey and Vine's (2011) series, which augments the gasoline price series (the solid line) with the time cost of waiting in gasoline lines (the dotted line). Without including the effects of rationing, the oil price changes in the 1970s look smaller than the later changes.

The third reason we should expect the effects to change over time is because of changes in the share of oil that is imported, as discussed in

Figure 2. The Gasoline Price and the Time Cost of Waiting in Gasoline Lines, 1967–2016[a]

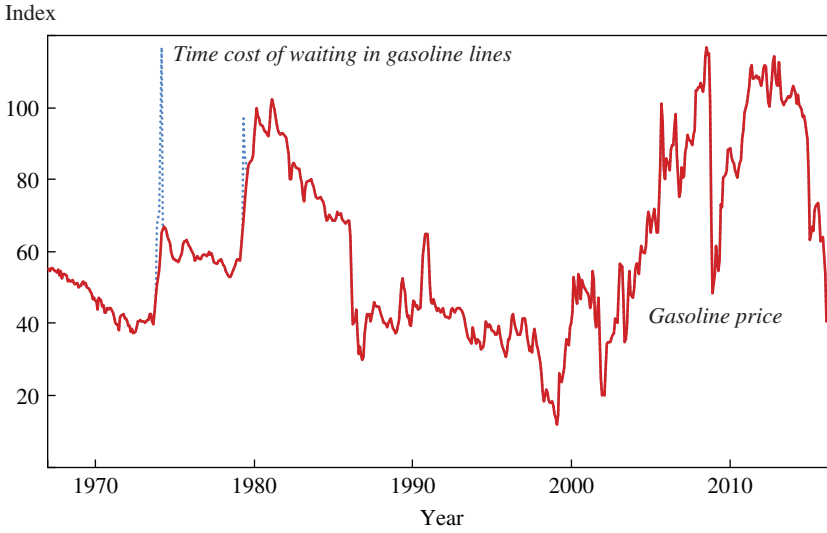

Sources: Ramey and Vine (2011); Federal Reserve Economic Data.
a. This figure is an extension of Ramey and Vine's (2011) figure 2. The base year is January 2011.

the simple theory presented above. The oil shale revolution has reduced the import share of petroleum significantly in recent years, as my figure 3 shows. This decline in the import share means that the term in parentheses in my equation 1 is now much lower. Because this term multiplies the relative price term, it implies that the impact of a change in real oil prices on other consumption should be lower now.

In sum, there are multiple reasons to expect oil price shocks to have smaller effects now than they did in the 1970s. Thus, I question Baumeister and Kilian's use of estimates that assume a constant effect since 1970.

CONCLUSIONS Baumeister and Kilian have presented a comprehensive analysis of the likely effects of the oil price decline during the past two years. Although I have raised questions about particular aspects of their analysis, it is reassuring that their conclusion is very similar to that of the 2016 *Economic Report of the President,* which focuses entirely on the terms-of-trade effects. That report concludes, "The decline in oil prices had the direct impact of boosting real GDP growth by 0.1 percentage point during 2014 and 0.2 percentage point during 2015" (Council of Economic Advisers 2016, p. 55). These numbers are very close to those estimated by Baumeister and Kilian, who find a small boost of about 0.2 percentage

Figure 3. Net Imports as a Share of the Petroleum Supply, 1950–2015

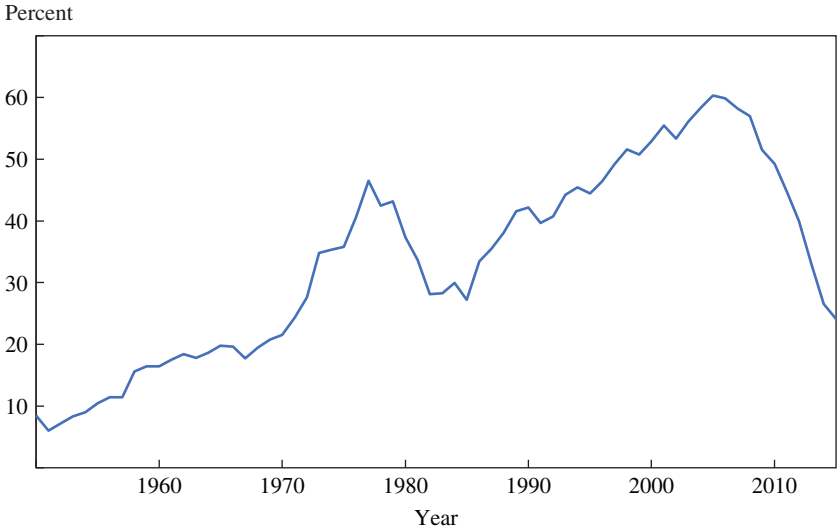

Source: U.S. Energy Information Administration (2016).

point per year, on average, from 2014 to 2016. Thus, it appears that there is a consensus that the recent fall in oil prices probably had little net effect.

REFERENCES FOR THE RAMEY COMMENT

Backus, David K., and Mario J. Crucini. 2000. "Oil Prices and the Terms of Trade." *Journal of International Economics* 50, no. 1: 185–213.

Blanchard, Olivier, and Jordi Galí. 2010. "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s so Different from the 1970s?" In *International Dimensions of Monetary Policy,* edited by Jordi Galí and Mark Gertler. University of Chicago Press.

Council of Economic Advisers. 2016. *Economic Report of the President.* Washington: U.S. Government Publishing Office.

Edelstein, Paul, and Lutz Kilian. 2009. "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics* 56, no. 6: 766–79.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95, no. 1: 161–82.

Ramey, Valerie A., and Daniel J. Vine. 2011. "Oil, Automobiles, and the U.S. Economy: How Much Have Things Really Changed?" *NBER Macroeconomics Annual* 25: 333–67.

U.S. Energy Information Administration. 2016. *Monthly Energy Review,* October. Washington: U.S. Department of Energy.

Yellen, Janet L. 2011. "Commodity Prices, the Economic Outlook, and Monetary Policy." Speech given at the Economic Club of New York, New York, April 11.

GENERAL DISCUSSION Jonathan Pingle wondered if the authors had considered looking at the Brent–WTI spread (which refers to the Brent crude oil spot price minus the West Texas Intermediate, or WTI, crude oil spot price). He noted that the Brent–WTI spread fluctuated during the authors' sample period, and that, historically, domestic U.S. gasoline prices have moved more closely with the Brent crude oil price, while the production of U.S. gasoline has a greater influence on the WTI price. He also noted that the U.S. Bureau of Labor Statistics' Consumer Expenditure Survey tentatively indicates there is evidence that the response of spending and consumption to an oil or energy price shock is lower for older age groups. Given the large shift in the U.S. population's age distribution since the 1970s, one might expect some differential consumption responses today compared with then.

Steven Braun had no criticism of the paper, but suggested a simpler way to think about the results, using back-of-the-envelope calculations. The United States imports about 1.8 billion barrels of oil per year on net (that is, oil coming in minus refined products going out). Therefore, a $10 increase in the price of a barrel of oil costs the U.S. economy $18 billion. Because nominal annual U.S. GDP is about $18 trillion, this $10 price increase translates into a 0.1 percent positive effect on the economy. In mid-2014, the price of oil was about $100 per barrel; and by the end of 2015, it was about $40 per barrel. Using the logic given here, this $60 difference should translate into about a 0.6 percent boost to the U.S. economy on the demand side. On the supply side, the national income accounts show that during this same period, the price decline led to about a 0.3 percent decrease in the contribution of oil and natural gas drilling investment to overall GDP. Therefore, the net effect of the oil price decline on U.S. GDP should be about 0.6 percent (from the demand side), minus 0.3 percent (from the supply side), which is 0.3 percent. This result follows the 2016 *Economic Report of the President*, and is close to the authors' estimate of the net stimulus reported in table 8.[1]

Steven Davis followed up on a comment made by discussant James Hamilton about the likelihood that specialized capital has been underutilized

1. Council of Economic Advisers, *Economic Report of the President* (2016): 55–57.