Is the Discretionary Income Effect of Oil Price Shocks a Hoax?

Christiane Baumeister
University of Notre Dame
CEPR

Lutz Kilian
University of Michigan
CEPR

Xiaoqing Zhou
Bank of Canada

June 7, 2017

The transmission of oil price shocks has been a question of central interest in macroeconomics since the 1970s. There has been renewed interest in this question after the large and persistent fall in the real price of oil in 2014-16. In the context of this debate, Ramey (2017) makes the striking claim that the existing literature on the transmission of oil price shocks is fundamentally confused about the question of how to quantify the effect of oil price shocks. In particular, she asserts that the discretionary income effect on private consumption, which plays a central role in contemporary accounts of the transmission of oil price shocks to the U.S. economy, makes no economic sense and has no economic foundation. Ramey suggests that the literature has too often confused the terms-of-trade effect with this discretionary income effect, and she makes the case that the effects of the oil price decline of 2014-16 on private consumption are smaller for a multitude of reasons than suggested by empirical models of the discretionary income effect. We review the main arguments in Ramey (2017) and show that none of her claims hold up to scrutiny. Our analysis highlights the theoretical basis of the discretionary income effect. We also discuss improved regression-based estimates of this effect that allow for changes in the dependence on oil and gasoline imports, and we highlight the fact that alternative estimates used by policymakers involve strong simplifying assumptions.

JEL codes: C51, Q43
Key Words: Stimulus; oil price decline; discretionary income effect; expenditure share; gasoline; net oil imports.

Acknowledgments: We thank Jim Hamilton and two anonymous referees for helpful comments. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Bank of Canada.

Corresponding author: Lutz Kilian, Department of Economics, 611 Tappan Street, University of Michigan, Ann Arbor, MI 48109-1220, USA. Email: lkilian@umich.edu.

Non-Technical Summary

The transmission of oil price shocks has been a question of central interest in macroeconomics since the 1970s. There has been renewed interest in this question after the large and persistent fall in the real price of oil in 2014-16. For example, Baumeister and Kilian (2017) provided evidence that the overall effect of this oil price decline on U.S. real GDP growth has been close to zero, consistent with the absence of an economic boom in the U.S. economy after June 2014.

In the context of this larger debate, Ramey (2017) made the striking claim that the existing literature on the transmission of oil price shocks is fundamentally confused about the question of how to quantify the effect of oil price shocks. In particular, she makes the case that the discretionary income effect on private consumption, which plays a central role in contemporary accounts of the transmission of oil price shocks to the U.S. economy such as Kilian (2008, 2014), Hamilton (2009, 2013), and Blanchard and Galí (2010), has no economic foundation. Ramey suggests that her views are shared by U.S. policymakers.

This article reviews the main arguments in Ramey (2017) and contrasts her reasoning with that in the existing literature. Our analysis demonstrates that the discretionary income effect on private consumption is closely related to the terms-of-trade effect of a change in the real price of oil on real domestic income. They key difference is that it explicitly allows for the fact that consumers purchase motor fuel rather than crude oil. Speeches by Fed officials show that this view is shared by central bankers. Likewise, the Council of Economic Advisers recently employed a simplified version of the model used by Baumeister and Kilian (2017). In addition, we show that state-of-the-art regression methods for quantifying the consumption stimulus of unexpectedly low oil prices are consistent with the implications of New Keynesian DSGE models. Finally, we demonstrate that there is no evidence that the regressions used to estimate the economic stimulus for 2014-16 suffer from structural breaks. We conclude that none of Ramey's claims holds up to scrutiny, and that there is no reason to rewrite the literature on the transmission of oil price shocks.

1. Introduction

The transmission of oil price shocks has been a question of central interest in macroeconomics since the 1970s.[1] There has been renewed interest in this question after the large and persistent fall in the real price of oil in 2014-16 (see, e.g., Yellen (2015), Council of Economic Advisers (2015, 2016), Baumeister and Kilian (2017)). In the context of this debate, Ramey (2017) makes the striking claim that the existing literature on the transmission of oil price shocks is fundamentally confused about the question of how to quantify the effect of oil price shocks. In particular, she makes the case that the discretionary income effect on private consumption, which plays a central role in contemporary accounts of the transmission of oil price shocks to the U.S. economy such as Kilian (2008, 2014), Hamilton (2009, 2013), and Blanchard and Galí (2010), "makes no economic sense" and "has no economic foundation". Ramey suggests that the literature "has too often confused the terms-of-trade effect … with this discretionary income effect" and makes the case that the effects of the oil price decline of 2014-16 on private consumption for a variety of reasons are smaller than estimated in Baumeister and Kilian (2017). Ramey also alleges that U.S. policymakers focus entirely on the terms-of-trade effect of changes in the real price of oil rather than the discretionary income effect.

Such strong statements raise the question of whether existing studies are really as confused as Ramey suggests or whether Ramey has perhaps misinterpreted the existing literature on the transmission of oil price shocks. Answering this question is of some importance and urgency, given the extent of recent oil price fluctuations. This article reviews the main arguments in Ramey (2017) and contrasts her reasoning with that in the existing literature. We show that none of Ramey's claims holds up to scrutiny. Our analysis demonstrates that the

[1] See, e.g., Fried and Schultz (1975), Sachs (1982), Hamilton (1988), Backus and Crucini (2002), Leduc and Sill (2004), Dhawan and Jeske (2008a), Edelstein and Kilian (2007, 2009).

discretionary income effect on private consumption is in fact closely related to the terms-of-trade effect of a change in the real price of oil on real domestic income and that U.S. policymakers are fully aware of this link. We conclude that there is no reason to rewrite the literature on the transmission of oil price shocks.

2. What is the Discretionary Income Effect?

Ramey argues that the literature has too often confused the terms-of-trade effect, which has a sound economic basis, with the discretionary income effect, which, according to Ramey, has no economic foundation. In particular, she asserts that the estimates of the discretionary income effect reported in Edelstein and Kilian (2009), Hamilton (2009, 2013), and Baumeister and Kilian (2017), among others, are invalid. It therefore is useful to start by briefly reviewing the nature of the discretionary income effect, as discussed in academic studies and policy documents.

The idea of the discretionary income effect was first articulated in the academic literature in Edelstein and Kilian (2009). In models of the discretionary income effect, an increase in the real price of imported crude oil (and hence in the real retail price of gasoline) is akin to a tax increase from the point of view of consumers, given that the demand for gasoline is price inelastic (see Coglianese et al. 2017). As this "oil tax" is transferred abroad, the aggregate income available for other purchases declines and private consumer spending falls, causing a decline in real GDP.[2] In response to a fall in the real price of oil, this discretionary income effect operates in reverse, and is expected to generate a stimulus for the U.S. economy.

Underlying this analysis is the view that real oil price shocks represent terms-of-trade

[2] Although some of the "oil tax" may ultimately be recycled, as oil-exporting countries directly or indirectly increase imports of U.S.-produced goods and services, this petrodollar recycling tends to occur with a considerable delay, if at all, and hence may be ignored in defining the shock to consumer's discretionary income.

shocks that alter domestic spending, which, through a Keynesian multiplier, affect real GDP growth. The main difference between Edelstein and Kilian's (2009) analysis and the conventional analysis of real oil price shocks as terms-of-trade shocks is that Edelstein and Kilian explicitly recognize the fact that consumers do not purchase crude oil, but rather refined oil products such as motor fuel. This fact requires a refinement of the standard terms-of-trade shock argument. Although only a small fraction of the gasoline consumed in the United States is imported, the domestic real price of gasoline responds directly to shocks to the real price of imported crude oil, to the extent that the gasoline sold in the United States is produced from imported crude oil. As the real price of gasoline changes in response to the higher cost of imported crude oil, so does the amount of real income that U.S. consumers have to give up to pay for the crude oil imports contained in the gasoline sold at gas stations.

Ramey (2017) denies the existence of this discretionary income effect on private consumption. First, she suggests that U.S. policymakers view oil price shocks as terms-of-trade shocks rather than shocks to consumers' discretionary income, implying that these explanations are mutually exclusive. Second, she asserts that the academic literature has incorrectly equated terms-of-trade shocks with relative price shocks. Specifically, she not only claims that existing studies ignored the fact that real oil price shocks are terms-of-trade shocks, but she also insists that the effect of real oil price shocks is invariant to the share of consumer spending on oil products. Third, Ramey claims that there is no theoretical support for the discretionary income effect. We address each of these claims in turn.

Ramey's first point about policymakers rejecting the notion of a discretionary income effect is clearly at odds with evidence from policy documents. For example, the Council of Economic Advisers in its *2015 Economic Report of the President* stresses that:

"the immediate effect of a price increase on an imported good like oil, which has a price-inelastic demand, is to decrease consumption of domestic goods and services, and, as a result, to decrease GDP." (p. 270)

The same report elaborates that:

"low oil prices benefit … the U.S. economy" because "lower fuel costs increase real household income and stimulate consumption … through lower gasoline prices" (p. 62).

The *2016 Economic Report of the President* is even more specific about this consumption response. It concludes that after June 2014 "lower gasoline prices freed up income for other purchases" (p. 54).

A second example is Yellen (2011) who observes that:

"… higher oil prices lower American income overall because the United States is a major oil importer and hence much of the proceeds are transferred abroad. […] Thus, an increase in the price of oil acts like a tax on U.S. households, and … tends to have a dampening effect on consumer spending. […] Staff analysis at the Federal Reserve Board indicates that a dollar increase in retail gasoline prices … reduces household disposable income ... and exerts a significant drag on consumer spending."

Contrary to the claim in Ramey (2017), these quotes contain every element of the discretionary income effect described by Edelstein and Kilian: (1) The transmission of oil price shocks to private consumption occurs through changes in the price of gasoline; (2) the demand for gasoline is price-inelastic; (3) the "oil tax" on consumer spending arises to the extent that the added spending on gasoline is transferred abroad; (4) this tax reduces discretionary income and dampens consumer spending, and, hence, real GDP.[3]

As noted by Baumeister and Kilian (2017), Yellen even acknowledges the fact that the effect of these gasoline price shocks depends on households' dependence on gasoline purchases, consistent with the standard regression models of the discretionary income effect which allow for changes in the gasoline expenditure share of consumers. To quote Yellen (2011):

[3] Discretionary income is sometimes referred to as disposable income. A likely reason is that the term-of-trade effect of higher oil prices may be viewed as an "oil tax" and disposable income is usually defined as after-tax income. Of course, this terminology is misleading because the "oil tax" is not considered a tax in national income accounting. More precisely, discretionary income is defined as disposable income minus real gasoline expenditures.

"… cheap oil encouraged households to purchase gas-guzzling cars …. Consequently, when oil prices quadrupled in 1973-74, that degree of energy dependence resulted in substantial adverse effects on real economic activity. Since then, however, energy efficiency in … consumption has improved markedly" and, as a result, the effect of oil price shocks on the real economy "has decreased substantially over the past several decades."

As to Ramey's second point, there is no evidence to support Ramey's claim that Edelstein and Kilian (2009) or related studies failed to interpret real energy price shocks as terms-of-trade shocks. In fact, Edelstein and Kilian (2009, p. 767) explicitly stressed that these relative price shocks take place in an open economy:

"Implicit in this view is the assertion that higher energy prices are primarily driven by higher prices for imported energy goods, and that at least some of the discretionary income lost from higher prices of imported energy goods is transferred abroad and is not recycled in the form of higher U.S. exports."

As these quotes show, Edelstein and Kilian (2009) and related studies agree with Yellen (2011) that shocks to the real price of imported commodities are terms-of-trade shocks. Edelstein and Kilian (2009) elaborate:

"In the case of a purely domestic energy price shock (such as a shock to U.S. refining capacity), it is less obvious that there is an effect on aggregate discretionary income. First, the transfer of income to the refiner may be partially returned to the same consumers in the form of higher wages or higher stock returns on domestic energy companies. Second, even if the transfer is not returned, higher energy prices simply constitute an income transfer from one consumer to another that cancels in the aggregate."

Thus, Ramey's assertion that Edelstein and Kilian (2009) confuse relative price shocks with terms-of-trade shocks is without basis. Ramey at some point of her discussion acknowledges this point, but suggests that the regression analysis in Edelstein and Kilian (2009) and related work remains flawed in that it proceeds, according to her, as if oil imports had nothing to do with this effect. This is not the case. In sections 3.1 and 3.2 of this paper, we examine in detail the regression models in question and make precise in what sense these models account for the role of oil imports. In fact, not only is there no clear separation between terms-of-trade shocks and discretionary income shocks, as shown in section 3, but Ramey's claim that the effect of these shocks is invariant to the share of consumer spending on oil products is mistaken. Moreover, we show that the refinements of the standard terms-of-trade shock argument proposed by Edelstein and Kilian (2009) are necessary to account for the fact that consumers purchase gasoline, as opposed to Ramey's (2017) assumption that consumers purchase crude oil.

Ramey (2017) suggests that the best way of illustrating why the discretionary income shock is distinct from a terms-of-trade shock is in the context of a back-of-the-envelope calculation provided in Baumeister and Kilian (2017). She proposes a counterexample intended to disprove the existence of the discretionary income effect. The point of Ramey's example is to show that a change in the relative price of a good subject to inelastic demand (say, the real price of services) has no direct effect on aggregate real consumption and real GDP. She concludes that therefore a change in the relative price of oil has no direct effect on real consumption and real GDP either. Closer examination, however, reveals that her example is not a counterexample to Edelstein and Kilian's (2009) analysis of the effect of changes in the real price of imported crude oil. Although the discretionary income of an individual consumer, all else equal, is reduced by the higher real price of services, if none of the services in question are imported, the income generated by purchases of services is recycled within the domestic economy, leaving aggregate consumer spending unaffected, as long as all consumers have the same propensity to consume.[4] Thus, the effect of a change in the real price of services on aggregate discretionary income and

[4] As discussed in Hamilton (2013, 2017), among others, to the extent that the marginal propensity to spend is not the same, an increase in the relative price may affect aggregate spending even in a closed economy. Given that the price elasticity of gasoline demand is comparatively low, an exogenous increase in the real price of gasoline causes a reduction in consumers' discretionary income. Although in a closed economy consumers' increased spending on gasoline represents income for someone else, by construction, it may take considerable time for this income to be returned to consumers in the form of company profits, royalties, or dividends paid to shareholders, or to be spent by oil companies in the form of increased investment expenditures. In a Keynesian model, differences in the marginal propensity to spend thus may affect the overall level of spending and hence the business cycle. Given our focus on the terms-of-trade effect, our discussion abstracts from this additional channel of transmission.

on aggregate spending is zero by construction, and there is no contradiction at all. Rather than undermining the argument in Baumeister and Kilian (2017), Ramey's example underscores the importance of carefully verifying all conditions for the existence of a discretionary income effect, as stated in Edelstein and Kilian (2009).

Finally, as to Ramey's claim that the discretionary income effect discussed in Edelstein and Kilian (2009) lacks any economic foundation, sections 3.4 and 3.5 show that this view is mistaken. The economic relationships underlying the regression models of the discretionary income effect can, in fact, be derived from a fully specified New Keynesian DSGE model, as described in Blanchard and Galí (2010). Our analysis shows that the discretionary income channel of the transmission of oil price shocks is equivalent to what Ramey refers to as a shift in domestic real income arising from shocks to the price of imported crude oil. This shift in real domestic income also is commonly referred to as an "oil tax", when discussing unexpected oil price increases (see, e.g., Bernanke (2006); Edelstein and Kilian (2009); Yellen (2011); Baumeister and Kilian (2017)). Bernanke (2006) and Edelstein and Kilian (2009), equivalently, referred to shifts in consumers' purchasing power associated with changes in the real price of imported crude oil. All these terms refer to essentially the same phenomenon.[5]

3. Are the Regressions Used in Estimating the Discretionary Income Effect Misspecified?

Baumeister and Kilian (2017) quantify the discretionary income effect based on estimates of a linear regression model of the relationship between changes in real U.S. private consumption and changes in consumers' purchasing power associated with gasoline price fluctuations, controlling

[5] Ramey herself at times refers to an "income effect". We do not find this terminology precise or useful, because the effect in question is distinct from the income and substitution effects associated with a relative price change in a closed economy. It only arises in response to an unexpected change in the price of an imported good whose demand is price inelastic. Likewise, this effect sometimes is called a wealth effect in applied work. We do not use this terminology because a wealth effect more properly refers to changes in the valuation of assets in response to oil price shocks (see, e.g., Kilian, Rebucci and Spatafora 2009).

for the evolution of the share of motor fuel expenditures in total consumer expenditures. Their approach is in turn a refinement of the original regression analysis in Edelstein and Kilian (2009). Ramey (2017) argues that this regression specification is, according to her, inconsistent with the interpretation of oil price shocks as terms-of-trade shocks, as discussed in Backus and Crucini (2000) and Kehoe and Ruhl (2008), for example, and fundamentally misspecified.

In making her case that the Baumeister-Kilian regression framework is misspecified, Ramey focuses on a simplified version of the standard national income accounting identity

$$Y = C + I + G + X - M,$$

where Y is real GDP, C is private consumption, I is investment, G is government spending, and $X - M$ is the external balance. Ramey postulates an economy, in which $I = G = 0$ and $X - M = 0$ such that

$$Y = C.$$

She then allows for two goods, oil and a non-oil good, such that in nominal terms

$$P_{oil}Y_{oil} + P_{noil}Y_{noil} = P_{oil}C_{oil} + P_{noil}C_{noil},$$

where P's denote prices, Y's denote quantities produced domestically and C's quantities consumed. Rearranging this identity yields

$$C_{noil} = Y_{noil} + \frac{P_{oil}}{P_{noil}}\left(Y_{oil} - C_{oil}\right). \quad (1)$$

Equation (1) highlights that if a country is a net oil importer, so $Y_{oil} - C_{oil} < 0$, an increase in the real price of oil will lower domestic income, whereas it will raise real income if $Y_{oil} - C_{oil} > 0$. There is, of course, universal agreement on this accounting identity and its interpretation. What is not clear is why Ramey (2017) concludes from this equation that the regression models used in Edelstein and Kilian (2009), Hamilton (2009), and Baumeister and Kilian (2017), among others, are fundamentally misspecified and make no economic sense. The remainder of this section seeks to clarify the merits of this claim.

Equation (1) suggests a tight relationship between non-oil consumption and net oil imports. It should be noted, however, that, in a more realistic model, that relationship is not clear-cut at all. If we allow for nonresidential investment expenditures in this accounting identity, for example, it becomes indeterminate by how much non-oil consumer expenditures and by how much business investment expenditures must fall to accommodate the loss in real income experienced by a net oil importing economy. This means that the argument outlined by Ramey is incomplete. We need to specify a mechanism to explain by how much non-oil consumer expenditures in this economy are affected by this terms-of-trade shock. Edelstein and Kilian (2009) propose such a mechanism, consisting of four elements:

1. As the real price of crude oil increases, so does the real price of gasoline. The extent of this price increase depends on the cost share of crude oil in producing gasoline.
2. Because the demand for gasoline is price-inelastic, consumers spend more on gasoline than before the gasoline price increase.
3. To the extent that the revenue from gasoline sales is transferred abroad and not returned to the U.S. economy, consumers' aggregate discretionary income (defined as after-tax real income minus real gasoline expenditures) falls, resulting in lower domestic aggregate demand.
4. In a Keynesian model, this reduction in aggregate demand may cause a decline in real GDP.

It is important to stress that the extent to which discretionary income falls depends on the extent to which the economy relies on net imports of crude oil. If the economy were self-sufficient in crude oil, for example, there would be no change in consumers' aggregate discretionary income. Ramey (2017) suggests that the mechanism described by Edelstein and Kilian (2009) implies that any fall in relative prices would cause a consumption stimulus, as long as demand is inelastic. Her interpretation, however, is missing Edelstein and Kilian's point that a discretionary income effect arises only to the extent that the real gasoline expenditures are transferred abroad and are not returned to the U.S. economy in the form of higher U.S. exports.

Next, we elaborate on a number of special cases of equation (1). It is useful to start the discussion with the case of a country that does not produce any crude oil domestically. All oil is imported. The situation of a country which produces crude oil domestically in addition to importing crude oil is discussed in sections 3.2 and 3.3.

3.1. Case 1: The Baseline Model without Domestic Oil Production

In the absence of domestic oil production, equation (1) simplifies to

$$C_{noil} = Y_{noil} - \frac{P_{oil}}{P_{noil}} C_{oil}, \qquad (1')$$

where P_{oil}/P_{noil} may also be viewed as the terms of trade because oil is the import good and the non-oil good is the export good. In other words, an increase in the price of oil expressed in units of the non-oil consumption good, all else equal, reduces the amount of the non-oil consumption good available for domestic consumption, because more of this good must be used to pay for the crude oil imports. Equivalently, we may write equation $(1')$ as $Y_{noil} - C_{noil} = P_{oil}C_{oil}/P_{noil}$, highlighting the fact that crude oil imports must be financed by a trade surplus in the non-oil consumption good. This case is typical for many oil-importing countries. It also provides a useful starting point for our discussion.

Although equation (1′) helps us understand the central tradeoff between domestic consumption and crude oil imports, it is overly simplistic, even if we follow Ramey in abstracting from investment expenditures. Most importantly, this accounting identity is misspecified in that crude oil is not part of personal consumption expenditures. Consumers purchase refined products such as heating oil, diesel, and gasoline rather than crude oil. This distinction matters. Given the cost share of crude oil in producing gasoline of about half, a 50% decline in the real price of crude oil translates to a 25% decline in the retail price of gasoline only. Hence, it does not make economic sense to relate consumer purchase decisions to changes in the real price of crude oil, as proposed by Ramey (2017). Instead, we need to focus on consumer expenditures on refined products and the corresponding product prices. Of these products, heating oil may be safely ignored because most U.S. households rely on other forms of home heating including natural gas and electricity (see Davis and Kilian 2011). Hence, Baumeister and Kilian (2017), for example, focus on consumer expenditures on motor fuel, which we will refer to, somewhat imprecisely, as gasoline in this paper.

A key point of contention is Ramey's (2017) suggestion that – in the context of case 1 – we should treat the increase in the real value of the imported oil (or its gasoline equivalent), caused by an unexpected increase in the real price of oil, as the terms-of-trade shock in regression analysis. She proposes relating the growth rate of overall consumption to this measure of the terms-of-trade shock. This argument is not persuasive. Clearly, the negative stimulus associated with this unexpected decline in consumer purchasing power will be larger in an economy in which most of domestic output is used to pay for energy imports than in an economy that is not very dependent on energy imports. Put differently, if a given increase in real domestic income is enough to raise the consumption growth rate by 2 percent, when gasoline consumption amounts to 50% of total consumption, one would not expect the same stimulus of 2 percent, when gasoline consumption only amounts to 1% of total consumption. Thus, it is essential that we normalize the real gasoline price shock, before fitting regression models. A formal derivation of this point is discussed in section 3.5.

Based on the work of Edelstein and Kilian (2009), we can address this concern by measuring the increase in consumer purchasing power arising from lower gasoline prices by

$$-\frac{C_t^{gas} P_{t+1}^{gas} / P_{t+1}^{PCE} - C_t^{gas} P_t^{gas} / P_t^{PCE}}{C_t^{ngas} P_t^{ngas} / P_t^{PCE}}, \tag{2}$$

where $C_t^{gas} P_t^{gas} / P_t^{PCE}$ denotes the real cost of purchasing C_t^{gas} in period t at the old gasoline price P_t^{gas}, $C_t^{gas} P_{t+1}^{gas} / P_{t+1}^{PCE}$ denotes the real cost of purchasing C_t^{gas} in the following period at the new gasoline price P_{t+1}^{gas}, and $C_t^{ngas} P_t^{ngas} / P_t^{PCE}$ denotes the real cost of purchasing the non-gasoline consumption good, all expressed in units of the domestic consumer basket, valued at P_t^{PCE}.[6] The central idea of Edelstein and Kilian is that we need to focus on the change in the real price of gasoline triggered by the shock to the real price of imported crude oil in quantifying the effects of this terms-of-trade shock on private consumption. Equation (2) expresses the change in real expenditures on oil imports (expressed as their gasoline equivalent) as a fraction of real expenditures on the non-gasoline consumer good. Multiplying this expression by P_t^{gas} / P_t^{gas} and rearranging yields

$$-\left(\frac{C_t^{gas} P_t^{gas}}{C_t^{ngas} P_t^{ngas}}\right) \cdot \left(\frac{P_{t+1}^{gas} / P_{t+1}^{PCE} - P_t^{gas} / P_t^{PCE}}{P_t^{gas} / P_t^{PCE}}\right),$$

which can be approximated by

[6] Note that Edelstein and Kilian deflate all nominal prices by the overall consumer price index rather than the price index of the non-gasoline consumption good. Given the small expenditure share of gasoline, this difference is immaterial in practice.

$$-\left(\frac{C_t^{gas}P_t^{gas}}{C_tP_t^{PCE}}\right)\cdot\left(\frac{P_{t+1}^{gas}/P_{t+1}^{PCE}-P_t^{gas}/P_t^{PCE}}{P_t^{gas}/P_t^{PCE}}\right),$$

given that the share of gasoline in total consumer expenditures is small. In other words, the change in consumer purchasing power (or discretionary income), denoted by PP_t, that is associated with the terms-of-trade shock approximately equals the percent change in the real price of gasoline weighted by the nominal share of gasoline expenditures in total consumer expenditures.[7] By construction, $PP_t < 0$, if the real price of gasoline rises unexpectedly, and $PP_t > 0$, if the real price of gasoline falls unexpectedly. Under the additional empirically supported assumption that consumers employ a no-change forecast of the real price of gasoline, PP_t may be viewed as a shock to discretionary income or, equivalently, to consumers' purchasing power (see, e.g., Anderson et al. (2011); Baumeister and Kilian (2016, 2017)). Given the evidence in Kilian and Vega (2011), we impose the identifying assumption that PP_t is predetermined with respect to the U.S. economy. As shown in Baumeister and Kilian (2017), this allows one to consistently estimate the cumulative effects of purchasing power shocks from the stationary regression model

$$\Delta c_t = \sum_{i=1}^{6}\beta_i\Delta c_{t-i} + \sum_{i=0}^{6}\gamma_i PP_{t-i} + u_t, \quad (3)$$

where u_t denotes the regression error and Δc_t the percentage growth rate of U.S. real private consumption. Thus, far from ignoring the role of oil imports and gasoline imports, this specification is consistent with a world in which all oil is imported. This is not the only interpretation of this regression model, however, as shown next.

---

[7] In applied work, it is common to approximate the gasoline expenditure share not by the share at the end of the month preceding the real gasoline price change, but rather by the average gasoline expenditure share in the current month.

3.2. Case 2: A Model with a Constant Share of Domestic Oil Production in Oil Consumption

Having discussed the special case of an economy without domestic oil production, we now turn to the more general case in which some crude oil is produced domestically and some is imported. Recall that in this case

$$C_{noil} = Y_{noil} + \frac{P_{oil}}{P_{noil}} \left(Y_{oil} - C_{oil} \right). \tag{1}$$

Equivalently, we may write that $Y_{noil} - C_{noil} = - P_{oil} \left(Y_{oil} - C_{oil} \right) / P_{noil}$ such that net oil imports must be financed by a trade surplus in the non-oil good. In the case of the United States, $C_{oil} > Y_{oil}$ and the oil trade balance is negative. It is useful to first consider the case of a constant oil import share. For expository purposes, suppose that every period, half of the crude oil consumed by the economy is produced domestically and the other half is imported. In that case $Y_{oil} = 0.5 C_{oil}$, and we can write

$$C_{noil} = Y_{noil} - \frac{P_{oil}}{P_{noil}} 0.5 C_{oil}. \tag{1''}$$

It may seem that regression model (3) would be invalidated by this change in assumptions, but this is not the case. The key difference, after expressing this model equivalently in terms of gasoline prices and quantities, is that the change in purchasing power now only depends on 50% of real gasoline consumption:

$$-\frac{0.5 C_t^{gas} P_{t+1}^{gas} / P_{t+1}^{PCE} - 0.5 C_t^{gas} P_t^{gas} / P_t^{PCE}}{C_t^{ngas} P_t^{ngas} / P_t^{PCE}}., \tag{2''}$$

As a result, the change in purchasing power now is only half as large as before:

$$-0.5 \left(\frac{C_t^{gas} P_t^{gas}}{C_t P_t^{PCE}} \right) \cdot \left(\frac{P_{t+1}^{gas} / P_{t+1}^{PCE} - P_t^{gas} / P_t^{PCE}}{P_t^{gas} / P_t^{PCE}} \right).$$

This change in the definition of PP_t does not affect the estimate of the cumulative effect of the purchasing power shocks based on model (3), however, because the factor of 0.5 is automatically absorbed in the estimate of the slope coefficient, when fitting the original model specification (3) without adjusting the PP_t measure. Put differently, as long as the share of the oil consumed by the U.S. economy is constant, the estimates of the cumulative effect of the terms-of-trade shock on private consumption remain valid. We do not even need to know what that share is. Thus, the implicit assumption in the work of Edelstein and Kilian (2009) and others was not that there is no domestic oil production, but that the share of domestic oil production is approximately constant. The model is much more general than it may have seemed at first sight and allows for a share of oil imports in oil consumption of less than 1.

3.3. Case 3: A Model with a Variable Share of Domestic Oil Production in Oil Consumption

Of course, the dependence of the U.S. economy on imports of crude oil and gasoline has not remained constant over time, especially not in recent years during the shale oil revolution (see Kilian 2017). Figure 1 plots the fraction of U.S. consumer expenditures on gasoline that is transferred abroad. The rationale for this measure is as follows. If all gasoline consumed in the United States were imported, there would be no need for an adjustment of consumer expenditures and the weight would be 1. If some gasoline is imported, the proceeds from gasoline sales going directly abroad can be captured by weighting the gasoline expenditure share by the share of gasoline that is imported. We also need to incorporate the share of gasoline that is not imported, however, to the extent that this gasoline is produced from imported crude oil. That component is captured by weighting the fraction of gasoline produced domestically by the fraction of net crude oil imports in the total use of crude oil by the U.S. economy. Figure 1 shows the sum of these two components, expressed as a percentage of consumer expenditures on gasoline.

Figure 1 shows that this import dependence measure generally increases, as the price of oil rises, and decreases, as the price of oil falls. The average share is 49%, but at times the share dropped to 25% or reached 70%. To the extent that this share at a given point in time is above or below the average share of 49%, one would expect the regression model (3) to overpredict or underpredict the cumulative effects of terms-of-trade shocks associated with oil price fluctuations. This point has indeed been ignored in the existing literature or has been thought to be of secondary importance. As a result of the fluctuations in this share, one would expect regression model (3) to be reasonably accurate only during periods when the share of domestic oil production in the use of oil is close to its long-run average. That was indeed the argument made by Baumeister and Kilian (2017) in defense of their original approach to modelling the 2014Q3-2016Q1 episode. That study observed that in mid-2014 the share of domestic oil production was close to its long-run average of 49%, suggesting that the estimates of model (3) should be useful in analyzing this episode. In contrast, for the episode of the 1986 oil price decline, this assumption is invalid because at that point in time the share was only 34% (see Figure 1).

One way of handling this concern would be the use of time-varying coefficient regression models, but such models are prone to overfitting and their estimates can be sensitive to the choice of Bayesian priors. As shown by Baumeister and Kilian (2017), an alternative and more direct way of quantifying the importance of changes in the dependence of the U.S. economy on oil and gasoline imports is to directly incorporate this evolution in the construction of the purchasing power shock. A simple approximation is to weight U.S. consumer expenditures on gasoline by the share of the proceeds going abroad, resulting in an alternative definition of purchasing power shocks,

$$-\frac{C_t^{gas} P_t^{gas}}{C_t P_t^{PCE}} \cdot \left(\frac{P_{t+1}^{gas} / P_{t+1}^{PCE} - P_t^{gas} / P_t^{PCE}}{P_t^{gas} / P_t^{PCE}} \right) \cdot \left(s_t^{gasoline\ imports} + (1 - s_t^{gasoline\ imports}) s_t^{net\ oil\ imports} \right),$$

where $s_t^{gasoline\ imports}$ is the seasonally adjusted share of U.S. motor gasoline imports in total U.S. motor gasoline consumption, as reported by the EIA, and where $s_t^{net\ oil\ imports}$ is the seasonally adjusted share of U.S. net crude oil imports in the total use of crude oil by the U.S. economy, as defined in Kilian (2017). The adjustment factor by which we multiply the original PP_t measure corresponds to the share shown in Figure 1. This alternative measure of shocks to consumer's purchasing power (or discretionary income), denoted by $PP_t^{alternative}$, is not only more relevant for understanding the dependence of U.S. consumers on imports of crude oil and gasoline than the share of net imports in petroleum products supplied, as reported by Ramey (2017), but it also avoids the ad hoc aggregation of crude oil and refined products. It should be kept in mind, however, that even $PP_t^{alternative}$ is only an approximation, because it ignores changes in oil and gasoline inventories, because it assumes that the net share of imported crude oil is the same in the production of all refined products, because it does not differentiate between gasoline and other motor fuel, and because it makes no allowance for changes over time in the extent of petrodollar recycling from abroad.

One advantage of this alternative specification is that we can directly evaluate the empirical content of Ramey's concern regarding changes in the dependence of the U.S. economy on crude oil imports by comparing the estimates of the discretionary income effect based on baseline specification (3) with the estimates based on the alternative specification

$$\Delta c_t = \sum_{i=1}^{6} \beta_i \Delta c_{t-i} + \sum_{i=0}^{6} \gamma_i PP_{t-i}^{alternative} + u_t, \qquad (3')$$

which controls for variation in the dependence of the U.S. economy on imports of crude oil and gasoline. We can also verify the conjecture by Baumeister and Kilian (2017) that we may still rely on the baseline model, as long as the import share is near its long-run average in the episode of interest.

Figure 2 compares the original and the alternative measure of purchasing power shocks. The purchasing power shock used in the baseline model (shown in the upper panel) without loss of generality has been scaled by the average adjustment factor of 0.49, so the magnitude of the purchasing power shocks can be compared directly to that of the shocks in the lower panel, which have been adjusted by the factor shown in Figure 1. The largest differences arise during 2005-09. At many other times, the differences are negligible. For example, Figure 2 shows that the purchasing power shock under the alternative specification is much smaller during 1986Q1-1987Q3. In contrast, during 2014Q3-2016Q1, the purchasing power shocks are remarkably similar under both specifications. Thus, to the extent that the estimate of the discretionary income effect during 2014Q3-2016Q1 is affected by the definition of the purchasing power shock, that difference must arise through differences in the slope parameter estimates of the regression model rather than the measurement of the purchasing power shocks.

The first column of Table 1 shows that, according to the baseline model (3), the decline in the price of oil after June 2014 raised U.S. real private consumption by 1.2 percentage points cumulatively by the end of 2016Q1. The implied stimulus for U.S. real GDP is 0.7 after accounting for the 69% share of private consumption in GDP and allowing for an import propensity of 15%. Controlling for the dependence on oil and gasoline imports, as shown in the second column of the table, implies a somewhat smaller stimulus of 0.9 percentage points for consumption (or 0.5 percentage points for real GDP).[8] Nevertheless, the estimates are in the same ballpark. In both cases, there is clear evidence of a modest stimulus associated with the discretionary income effect. As expected, controlling for the dependence on oil and gasoline imports changes the estimates without overturning the substantive conclusions based on model (3).

From a policy point of view, the difference between a stimulus of 0.7 percentage points for U.S. real GDP and a stimulus of 0.5 percentage points is negligible. Thus, Ramey's conjecture that the baseline results in Baumeister and Kilian (2017) are an artifact of their failure to model time variation in the dependence of the U.S. economy on imports of gasoline and crude oil is clearly rejected. We conclude that, on a priori grounds, the alternative specification $(3')$ proposed by Baumeister and Kilian (2017) may be more appealing, but that the two specifications generate similar results for the 2014-16 episode, contradicting Ramey's claim. Moreover, which specification is used does not affect the substantive conclusion in Baumeister and Kilian (2017) that the net stimulus for the U.S. economy implied by lower oil prices after June 2014 has been close to zero.

In the last two columns of Table 1 we repeat this exercise for the 1986-87 episode, which was characterized by a similar, if much smaller sustained decline in the real price of oil. For this episode, one would expect the baseline model to be less accurate, given that the economy was less dependent on oil and gasoline imports in 1985Q4 than in 2014Q2, as reflected in the alternative purchasing power measure in Figure 2. The baseline model (3) implies that, in the seven quarters following 1985Q4, U.S. real private consumption cumulatively increased by 0.8

[8] The estimates in Table 1 for the alternative model do not match those in Baumeister and Kilian (2017) exactly due to a data transcription error in that paper, which has been corrected in the current analysis. For example, the cumulative effect on consumption in model $(3')$ is 0.87 rather than 0.92, as reported in Baumeister and Kilian (2017).

percentage points, as a result of the decline in the real price of oil that started in early 1986, implying a stimulus of 0.5 percentage points for real GDP. In contrast, the alternative model $(3')$ suggests a cumulative increase in private consumption of only 0.4 percentage points and an implied stimulus for U.S. real GDP of only 0.2 percentage points (see Table 1). Although this difference is not much larger in absolute terms than in the 2014-16 episode, it is at least twice as large in relative terms. This example illustrates that the adjustment factor may matter, providing empirical support for Baumeister and Kilian's decision to model the 1986-87 episode based on model $(3')$ rather than (3).

3.4. Yet Another Definition of the Purchasing Power Shock?

Given Ramey's insistence that the approach described in sections 3.2 and 3.3 is at odds with the policy consensus, it is useful to relate our approach to the definition of purchasing power shocks used in Council of Economic Advisers (2014, p. 25):

$$\left[as_t + (1-a)\overline{s}\right] \cdot \left(\frac{P_{t+1}^{oil} - P_t^{oil}}{P_t^{oil}} \right),$$

where $s_t = (M_t^{oil} - X_t^{oil})P_t^{oil} \big/ (Y_t P_t^Y)$, $\overline{s}$ is the sample average of s_t, $(M_t^{oil} - X_t^{oil})P_t^{oil}$ denotes nominal net oil imports, Y_t is real GDP, and P_t^Y is the GDP deflator.[9] The parameter a is either 0 or 1. For $a = 0,$ the CEA measure of the purchasing power shock reduces to

$$\overline{s} \cdot \left(\frac{P_{t+1}^{oil} - P_t^{oil}}{P_t^{oil}} \right),$$

and for $a = 1,$ it reduces to

[9] To be precise, the Council of Economic Advisers (2014) works with net petroleum imports rather than net crude oil imports, where petroleum is defined as crude oil plus selected refined products, but they weight the nominal expenditure share for petroleum by the change in the price of crude oil. Thus, for expository purposes, it is preferable to treat the quantities of oil and gasoline as exchangeable and to denote net petroleum imports as net oil imports.

$$s_t \cdot \left(\frac{P_{t+1}^{oil} - P_t^{oil}}{P_t^{oil}} \right).$$

It can be shown that these definitions under suitable simplifying assumptions may be derived from the original and the alternative definition of the purchasing power shock, respectively, as discussed in sections 3.2 and 3.3. For expository purposes, we consider the case of $a = 1.$ Recall our alternative definition of the purchasing power shock as

$$-\frac{C_t^{gas} P_t^{gas}}{C_t P_t^{PCE}} \cdot \left(\frac{P_{t+1}^{gas} / P_{t+1}^{PCE} - P_t^{gas} / P_t^{PCE}}{P_t^{gas} / P_t^{PCE}} \right) \cdot \left(s_t^{gasoline\ imports} + (1 - s_t^{gasoline\ imports}) s_t^{net\ oil\ imports} \right),$$

where $s_t^{net\ oil\ imports} \equiv \left(M_t^{oil} - X_t^{oil} \right) P_t^{oil} \Big/ \left((Y_t^{oil} + M_t^{oil} - X_t^{oil}) P_t^{oil} \right),$ and suppose that $C_t = Y_t,$ which amounts to imposing that $I = G = X - M = 0,$ as in Ramey's analysis. Further suppose that gasoline and oil are the same good such that $s_t^{gasoline\ imports} = 0,$ $P_t^{gas} = P_t^{oil},$ and $C_t^{gas} = C_t^{oil},$ and suppose that there is no inflation, so the real and the nominal price of oil coincide. Finally suppose that the change in U.S. oil inventories is zero such that $C_t^{oil} = Y_t^{oil} + M_t^{oil} - X_t^{oil}$ (see Kilian 2017). Then simple substitution shows that the alternative purchasing power shock measure reduces to

$$-\frac{C_t^{oil} P_t^{oil}}{Y_t P_t^{Y}} \cdot \left(\frac{P_{t+1}^{oil} - P_t^{oil}}{P_t^{oil}} \right) \cdot \left(\frac{\left(M_t^{oil} - X_t^{oil} \right) P_t^{oil}}{C_t^{oil} P_t^{oil}} \right),$$

or, equivalently,

$$-\frac{\left(M_t^{oil} - X_t^{oil} \right) P_t^{oil}}{Y_t P_t^{Y}} \cdot \left(\frac{P_{t+1}^{oil} - P_t^{oil}}{P_t^{oil}} \right),$$

which up to the sign normalization is identical to the definition used by the Council of Economic Advisers (2014, 2016). Of course, there is no good reason for using any of these unnecessary simplifying assumptions, but this exercise illustrates that policymakers fully understand the need for normalizing by the expenditure share, directly refuting Ramey's (2017) claim in this regard. This discussion also helps explain why the approximate results reported in Council of Economic Advisers (2016), as discussed in section 5, are close to those in Baumeister and Kilian (2017). The reason is that both use fundamentally the same approach, except that the computations in Baumeister and Kilian (2017) are more precise.

3.5. What Is the Economic Rationale for Weighting Real Gasoline Price Shocks by the Gasoline Expenditure Share?

In subsection 3.1, we explained intuitively why shocks to the real price of gasoline must be weighted by the gasoline expenditure share of consumers, when quantifying the discretionary income effect. Ramey (2017) makes the case that this approach is not supported by economic theory. She insinuates that fully specified economic models of the transmission of oil price shocks imply that the response of private consumption does not depend on consumers' energy share, rendering standard regression estimates of the discretionary income effect invalid.

This claim is not correct. Even granting that none of the currently available DSGE models provides a satisfactory representation of the real world, as noted by Kilian (2014), and that their quantitative implications are sensitive to ad hoc modelling assumptions, the fact that the impact of energy price shocks on the economy depends on the share of firm and household energy expenditures is well established. In fact, declines in the share of energy in production and consumption have been held responsible by many researchers for the reduced overall effect of oil price shocks on the U.S. economy in the literature (see, e.g., Edelstein and Kilian (2009), Blanchard and Galí (2010), Kilian and Vigfusson (2017)). The same point was made by Yellen (2011), as discussed earlier.

It is useful to review the theoretical rationale for the approach described in the preceding subsections. The DSGE model of Backus and Crucini (2000) referred to by Ramey is not helpful in answering that question because in that model households do not directly consume energy. Oil only enters into the production function. That assumption is typical for much of the DSGE literature on the transmission of oil price shocks, but fails to capture the reality that consumer spending on motor fuel is a major channel of the transmission of oil price shocks (see, e.g., Leduc and Sill 2004). In more recent DSGE models such as Dhawan and Jeske (2008a,b), which explicitly allow for household energy consumption, however, the responses of output to an exogenous real energy price shock depend on the household energy share. The economic rationale for weighting real gasoline price shocks perhaps is best illustrated by the New Keynesian DSGE model described in Blanchard and Galí (2010), which allows oil to enter both the utility function and the production function. The model makes no distinction between crude oil and gasoline. Blanchard and Galí (p. 418) show that, in equilibrium, the relationship between real private consumption and real value added in an oil-importing economy is approximately

$$\log\left(C_t\right) = \log\left(Y_t\right) - \left(\frac{\alpha}{1-\alpha} + \frac{C_t^{oil} P_t^{oil}}{C_t\, P_t^{CPI}}\right)\log\left(\frac{P_t^{oil}}{P_t^{PPI}}\right),$$

where C_t is real private consumption, Y_t is real value added, α is the oil share in production, $C_t^{oil} P_t^{oil} \big/ \left(C_t\, P_t^{CPI}\right)$ is the oil expenditure share of consumers, and $P_t^{oil} \big/ P_t^{PPI}$ is the real price of oil. This expression may be viewed as a generalization of equation (1). Given that α is quite small and that there is little difference between the PPI and the CPI, this expression can be approximated by

$$\log\left(C_t\right) \approx \log\left(Y_t\right) - \left(\frac{C_t^{oil} P_t^{oil}}{C_t\, P_t^{CPI}}\right)\log\left(\frac{P_t^{oil}}{P_t^{CPI}}\right),$$

which provides direct theoretical support for regression specification (3) and refutes Ramey's claim that this specification has no theoretical support.

4. Are the Regression Estimates in Baumeister and Kilian (2017) Invalidated by Structural Breaks?

In addition to questioning the motivation of the regression model used to estimate the discretionary income effect and its specification, Ramey also raises concerns about the structural stability of these regressions. She presents evidence that the cumulative impulse response of private consumption changes, when splitting the sample in January 1993. It is worth noting that her evidence is not based on the regression specification of Baumeister and Kilian (2017). Rather she reports results of a local projection model for the log-level of private consumption. This model specification involves estimating many more parameters than the original specification in Baumeister and Kilian (2017) and is known to be unreliable in small samples, so the estimates reported by Ramey have to be viewed with some caution (see Kilian and Kim (2011); Kilian and Lütkepohl (2017)). Clearly, however, qualitatively similar facts could be established for the original specification in Baumeister and Kilian (2017). In fact, as Ramey (2017) concedes, there is nothing new about her evidence. Essentially the same point was already made in Edelstein and Kilian (2009) who also discussed the causes of this time variation.

In discussing the structural stability of the estimates in Baumeister and Kilian (2017) it is useful to work with the exact specification used in that paper. Table 2 reports the cumulative response of real private consumption to a one standard deviation purchasing power shock after 20 months (corresponding to the length of the 2014.7-2016.3 episode). All results are based on the baseline model (3). The first row shows the full-sample estimate implied by this model. The next two rows in the table report the corresponding estimates for the two subperiods considered by Ramey (2017). Table 2 shows that the estimate based on the subsample 1970.2-1992.12 is 3.4 times larger than the estimate based on the subsample 1993.1-2016.3.

The difference across subsamples is not nearly as large as implied by the inefficient local projection estimator employed in Ramey (2017), but still large. It is also irrelevant for Ramey's argument that the estimates in Baumeister and Kilian (2017) are misleading. If there were a structural change, as asserted by Ramey, the analysis in Baumeister and Kilian (2017) would be in error to the extent that the full-sample estimate (which Ramey did not report) differs from the estimate based on the second subsample. The value of the response estimate for the first subsample is irrelevant for that question. Table 2 shows that the full-sample estimate of 0.12 is not much larger than the estimate of 0.07 for the second subsample. It is this much more modest difference in estimates that requires an economic explanation.

This difference may still seem too large for comfort, but what Ramey fails to mention is that there are well-known reasons for the apparent instability in the regression coefficients, when the model is evaluated on subsamples. Such instability is in fact expected based on previous research, even in the absence of any structural breaks in the data generating process. Although the average responses of real consumption to purchasing power shocks may be estimated reliably using long enough samples, when considering short subsamples, these responses will change in magnitude and even in sign, as the composition of oil demand and oil supply shocks evolves over time, giving the mistaken appearance of structural instability, even when there is no structural change at all (see, e.g., Kilian (2008; 2009a,b); Kilian and Park (2009)). For example, as discussed in Edelstein and Kilian (2009), the real gasoline price in the second half of Ramey's sample is dominated by unexpected shifts in the flow demand for crude oil, which cushion the direct effect of oil price fluctuations on the U.S. economy. Thus, there is no mystery why the response estimates are much lower in the second half of the sample. This result simply reflects changes in the composition of oil demand and oil supply shocks.

This point may be illustrated empirically. The second subsample chosen by Ramey is known to include few important supply and storage demand shocks. Instead, it is dominated by flow demand shocks for oil, making it unrepresentative for the sample as a whole as well as for the period since June 2014 (see, e.g., Kilian and Murphy 2014). By slightly extending the second subsample to include the invasion of Kuwait in 1990, which constitutes an important shock to oil supply and to the storage demand for oil, the estimate for the second subsample increases from 0.0692 to 0.0997. The latter estimate is very close to the full-sample estimate of 0.1194 (see Table 2), suggesting that there is no instability at all. In fact, using model estimates based on the full sample, as in Baumeister and Kilian (2017), makes sense in analyzing the effects of the 2014-16 oil price decline, because the real price of oil during that episode was driven by a combination of different demand and supply shocks rather than primarily by flow demand shocks (see Baumeister and Kilian (2016); Kilian (2017)).

It is conceivable, of course, that there are other reasons for the decline in the response of private consumption that are unrelated to changes in the composition of oil demand and oil supply shocks. Ramey provides three specific reasons why she believes that the effect of oil and gasoline price shocks should have declined. First, she suggests that real oil price shocks caused larger overall terms-of-trade fluctuations in the 1970s and early 1980s than in recent years. This argument is missing the point. First, Ramey misrepresents Backus and Crucini’s work. There are no exogenous real oil price shocks in their model, so Backus and Crucini (2000) could not possibly have quantified the causal effects of real oil price shocks. What they actually showed was that the time variation in the dynamic correlation between the real price of oil and the U.S.

terms of trade reflects changes in the composition of oil demand and oil supply shocks. In their words:

"… since our dynamic general equilibrium model predicts that the economy responds differently to oil supply shocks than to other shocks, changes in their relative importance help to account for the unstable correlations in the data." (p. 185)

This point and its implications have already been addressed in our earlier discussion. Second, there is no a priori presumption that the real price of oil and the overall terms of trade of the U.S. economy should move proportionately. Clearly, the terms of trade are subject to many more shocks than real oil price shocks, so the weakening unconditional co-movement between the real price of oil and the overall U.S. terms of trade, as displayed in Figure 1 of Ramey (2017), is not evidence of a structural change in the data generating process.

Second, Ramey insists that the terms of trade are affected by U.S. gasoline price controls and rationing. She proposes scaling the nominal gasoline price underlying the regressions in Baumeister and Kilian (2017) by a multiplicative factor as used in Ramey and Vine (2011). This adjustment is intended to capture the cost to consumers of waiting at gas stations that arose during 1973.12-1974.5 and in 1979.5-1979.7, when the government imposed gasoline price ceilings in response to oil price increases. Because this waiting cost is not associated with a transfer of real income abroad, however, this adjustment must not be used in quantifying the discretionary income effect. This misunderstanding is closely related to Ramey's failure to appreciate that purchasing power shocks reflect real income transfers abroad.

Finally, Ramey argues that declines in the share of oil imports explain the smaller cumulative response of private consumption in the sample starting in January 1993. That argument was already dispensed with in section 3 where we showed that the response of private consumption in 2014-16 is quite similar, whether we allow for changes in the dependence on oil and gasoline imports or not. Likewise, Baumeister and Kilian (2017) stressed that their substantive conclusions are unaffected by this adjustment. Thus, the evidence and arguments in Ramey (2017) regarding potential structural instabilities in the regression model of Baumeister and Kilian (2017) are missing the point.

5. Whither regression analysis?

As stressed by Ramey (2017), the estimates in Baumeister and Kilian (2017) are very close to the "direct effects" of lower oil prices on U.S. real GDP growth reported in the *2016 Economic Report to the President* (Box 2-1, p. 55-58) based on a simple back-of-the-envelope calculation. This report argued that if one values all U.S. net imports of crude oil and petroleum products at the nominal price of crude oil, then, given the cumulative decline in this price since June 2014, all else equal, we would expect a 0.1% increase in U.S. real GDP in 2014 and an additional increase of 0.2% in 2015. The implied cumulative increase in real GDP after mid-2014 is about 0.3%, which is remarkably close to the net stimulus of 0.39% estimated in Baumeister and Kilian (2017). Likewise, for private consumption the estimate of 0.6 in the *Report* is close to the estimate of 0.7 in Baumeister and Kilian.

The overall tenor of Ramey's analysis is that the effects of unexpectedly lower oil prices on consumption should be systematically smaller than the estimates reported in Baumeister and Kilian (2017), yet at the same time she is forced to acknowledge that these estimates are similar to those obtained in the *Report*, which she considers correctly executed. Our analysis suggests that it is not an accident that the *Report* reaches conclusions similar to Baumeister and Kilian (2017). As shown in section 3.4, both studies measure purchasing power shocks in fundamentally the same way, the only difference being that Baumeister and Kilian (2017) dispense with the restrictive assumptions underlying the analysis in the *Report*.

This is not the only difference, however. The *Report* is careful to stress that care is required in interpreting its estimate of the “direct effect”, because “consumers for whom lower gasoline prices freed up income for other purchases … may take time … to make additional purchases, so the timing of the additional spending may lag the declines in oil prices” (p. 56). Nor do the estimates in the *Report* account for additional multiplier effects from changes in spending (p. 56). Thus, these direct estimates conceptually differ from the cumulative effect estimated by Baumeister and Kilian (2017). Unlike the estimates cited in the *Report*, Baumeister and Kilian’s (2017) estimates take account of typical delays in spending as well as multiplier effects.

Perhaps most importantly, the analysis in Baumeister and Kilian (2017) explicitly recognizes the fact that consumer responses depend on the real price of gasoline rather than on the real price of oil, and it takes account of the timing and persistence of changes in the real price of gasoline on a month-by-month basis, allowing more accurate estimates of the cumulative effects and providing more information about the evolution of these cumulative effects month by month. Their analysis also dispenses with a number of other simplifying assumptions, as discussed in section 3.4. Thus, the fact that the back-of-the-envelope estimates in the *Report* are in the same ball-park as those in Baumeister and Kilian (2017) was by no means obvious ex ante. Back-of-the-envelope computations of this type are best viewed as a preliminary approximation to be validated by more formal regression-based methods of the type discussed in this paper rather than as a substitute for more formal analysis.

6. Conclusion

The discretionary income effect is a central element of the transmission of oil price shocks, as discussed in the existing literature. Recently, Ramey (2017) suggested that empirical studies quantifying this effect such as Edelstein and Kilian (2009), Hamilton (2009, 2013), and Baumeister and Kilian (2017), among others, have no economic foundation. She not only disputed the very existence of a discretionary income effect, but she called into question the stability of commonly used regression models in this literature, and she even questioned the data used by these models. Ramey's critique is remarkable in that she does not provide any constructive alternative modeling approaches, but implies that regression models are simply not necessary to assess the questions of interest.

Our analysis showed that Ramey's central claim that the discretionary income channel makes no economic sense is mistaken. In fact, the shocks to discretionary income (also referred to as the purchasing power shocks in the literature) discussed in the literature are substantively identical with the shifts in real domestic income associated with a terms-of-trade shock stressed by Ramey. The specification of conventional regression models of the discretionary income effect is fully consistent with recent New Keynesian DSGE models of the transmission of oil price shocks. Moreover, these regression models are already designed to deal with many of the concerns raised by Ramey. We discussed generalizations of these models and assessed the robustness of standard specifications to the concerns raised by Ramey.

We also demonstrated that the model specifications and diagnostics used by Ramey to establish the structural instability of regression models of the transmission of oil shocks are flawed. We showed that her modelling approach ignores recent developments in the literature, and we explained why the data modifications proposed by Ramey are inappropriate in the context of this literature. Finally, we discussed the importance of regression analysis in quantifying the effects of oil price shocks. We concluded that there is no reason to rewrite the literature on the transmission of oil price shocks.

References:

Anderson, S.T., R. Kellogg, and J.M. Sallee (2013), "What Do Consumers Believe about Future Gasoline Prices," *Journal of Environmental Economics and Management*, 66, 383-403.

Backus, D.K., and M.J. Crucini (2000), "Oil Prices and the Terms of Trade," *Journal of International Economics,* 50(1), 185-213.

Baumeister, C., and L. Kilian (2016), "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us," *Journal of Economic Perspectives*, 30, 139-160, Winter.

Baumeister and Kilian (2017), "Lower Oil Prices and the U.S. Economy: Is This Time Different?", *Brookings Papers on Economic Activity*, Fall, 287-336.

Bernanke, B.S. (2006), "The Economic Outlook," Remarks before the National Italian American Foundation, New York, November 28.

Blanchard, O.J., and J. Galí (2010), "The macroeconomic effects of oil price shocks: Why Are the 2000s so different from the 1970s?" in: Galí, J., and M. Gertler (eds.), *International Dimensions of Monetary Policy*, University of Chicago Press: Chicago, IL, 373-421.

Davis, L.W., and L. Kilian (2011), "The Allocative Cost of Price Ceilings in the U.S. Residential Market for Natural Gas," *Journal of Political Economy*, 119(2), 212-241.

Coglianese, J, Davis, L.W., Kilian, L., and J.H. Stock (2017), "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand," *Journal of Applied Econometrics*, 32(1), 1-15.

Council of Economic Advisers (2014), "The All-Of-The-Above Energy Strategy as a Path to Sustainable Economic Growth," The White House, Office of the Press Secretary, Washington, DC.

Council of Economic Advisers (2015), *Economic Report of the President*, Government Printing Office, Washington, DC.

Council of Economic Advisers (2016), *Economic Report of the President*, Government Printing Office, Washington, DC.

Dhawan, R. and K. Jeske (2008a), "Energy Price Shocks and the Macroeconomy: The Role of Consumer Durables," *Journal of Money, Credit, and Banking*, 40, 1357-1377.

Dhawan, R. and K. Jeske (2008b), "What determines the Output Drop after an Energy Price Increase? Household or Firm Energy Share?" *Economics Letters*, 107, 202-205.

Edelstein, P., and L. Kilian (2007), "The Response of Business Fixed Investment to Energy Price Changes: A Test of Some Hypotheses about the Transmission of Energy Price Shocks," *B.E. Journal of Macroeconomics*, 7.

Edelstein, P., and L. Kilian (2009), "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics*, 56, 766-779.

Fried, E., and C. Schultz (1975), *Higher Oil Prices and the World Economy*, Brookings Institution, Washington, DC.

Hamilton, J.D. (1988), "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy*, 96, 593-617.

Hamilton, J.D. (2009), "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, 1, Spring, 215-259.

Hamilton, J.D. (2013), "Oil Prices, Exhaustible Resources, and Economic Growth," in Fouquet, R. (ed.), *Handbook of Energy and Climate Change*, Cheltenham, U.K.: Edward Elgar Publishing, 29-57.

Hamilton, J.D. (2017), Comments on 'Lower Oil Prices and the U.S. Economy: Is This Time Different?" by Christiane Baumeister and Lutz Kilian', *Brookings Papers on Economic Activity*, Fall, 337-343.

Kehoe, T.J., and K.J. Ruhl (2008), "Are Shocks to the Terms of Trade Shocks to Productivity?" *Review of Economic Dynamics*, 11, 804-829.

Kilian, L. (2008), "The Economic Effects of Energy Price Shocks," *Journal of Economic Literature*, 46(4), 871-909.

Kilian, L. (2009a), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review*, 99(3), 1053-1069.

Kilian, L. (2009b), "Comment on 'Causes and Consequences of the Oil Shock of 2007-08' by James D. Hamilton," *Brookings Papers on Economic Activity*, 1, Spring, 267-278.

Kilian, L. (2014), "Oil Price Shocks: Causes and Consequences," *Annual Review of Resource Economics*, 6, 133-154.

Kilian, L. (2017), "The Impact of the Fracking Boom on Arab Oil Producers," forthcoming: *Energy Journal*.

Kilian, L., and Y.J. Kim (2011), "How Reliable Are Local Projection Estimators of Impulse Responses?" *Review of Economics and Statistics*, 93(4), 1460-1466.

Kilian, L., and D.P. Murphy (2014), "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics*, 29, 454-478.

Kilian, L., and H. Lütkepohl (2017), *Structural Vector Autoregressive Analysis*, Cambridge

University Press, forthcoming.

Kilian, L, and C. Park (2009), "The Impact of Oil Price Shocks on the U.S. Stock Market," *International Economic Review*, 50, 1267-1287.

Kilian, L., Rebucci, A. and N. Spatafora (2009), "Oil Shocks and External Balances," *Journal of International Economics*, 77(2), 181-194.

Kilian, L., and C. Vega (2011), "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices," *Review of Economics and Statistics*, 93, 660-671.

Kilian, L., and R.J. Vigfusson (2017), "The Role of Oil Price Shocks in Causing U.S. Recessions," forthcoming: *Journal of Money, Credit, and Banking*.

Leduc, S, and K. Sill (2004), "A Quantitative Analysis of Oil Price Shocks, Systematic Monetary Policy, and Economic Downturns," *Journal of Monetary Economics*, 51, 781-808.

Ramey, V.A. (2017), Comments on 'Lower Oil Prices and the U.S. Economy: Is This Time Different?" by Christiane Baumeister and Lutz Kilian', *Brookings Papers on Economic Activity*, Fall, 343-351.

Ramey, V.A., and D.J. Vine (2011), "Oil, automobiles, and the US economy: How Much Have Things Really Changed?" *NBER Macroeconomics Annual*, 25, ed. Acemoglu, D., and M. Woodford,.Chicago: University of Chicago Press, 333-368.

Sachs, J.D. (1982), "The Oil Shocks and Macroeconomic Adjustment in the United States," European Economic Review, 18, 243-248.

Yellen, J.L. (2011), "Commodity Prices, the Economic Outlook, and Monetary Policy," speech at the Economic Club of New York, New York, NY, April 11.

Yellen, J.L. (2015), *Monetary Policy Report*, Board of Governors of the Federal Reserve System.

Table 1: Predicted Cumulative Stimulus from Lower Oil Prices on the U.S. Economy

	2014.Q3-2016.Q1		1986.Q1-1987.Q3	
	Baseline Model (3)	Alternative Model (3′)	Baseline Model (3)	Alternative Model (3′)
Effect of Lower Oil Prices on U.S. Real Private Consumption (%)	1.20	0.87	0.84	0.43
Implied Effect on U.S. Real GDP (%)	0.70	0.51	0.45	0.23

Notes: Models (3) and (3′) are described in the text. The alternative model differs from the baseline model in that it allows for variation in the dependence of the U.S. economy on imports of crude oil and gasoline, as shown in Figure 1.

Table 2: Cumulative Response of U.S. Consumption in Month 20 to a One Standard Deviation Purchasing Power Shock

Sample	Cumulative Response (Percent)
1970.2-2016.3 (full sample)	0.1194
1970.2-1992.12 (first subsample)	0.2380
1993.1-2016.3 (second subsample)	0.0692
1990.1-2016.3 (including invasion of Kuwait)	0.0997

Notes: All results are based on the baseline model (3).

Figure 1: Adjustment Factor for the U.S. Dependence on Gasoline and Crude Oil Imports, 1973.1-2016.3

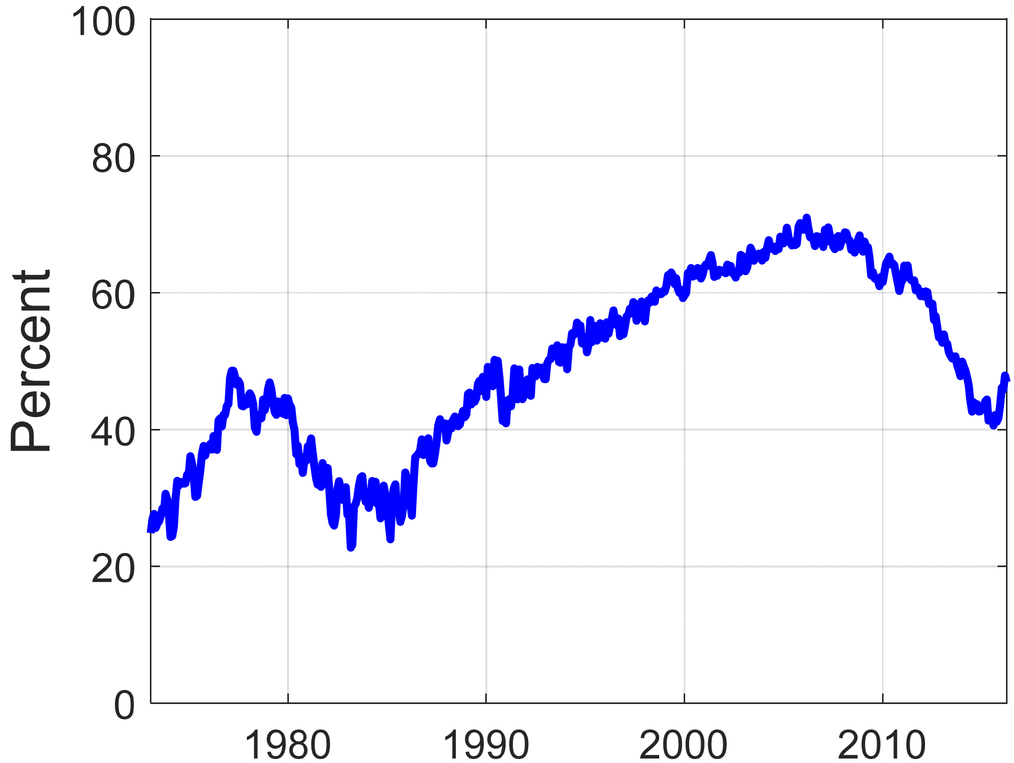

Notes: Authors' calculations based on data in the U.S. Energy Information Administration's *Monthly Energy Review*, as discussed in text. This adjustment factor measures the extent to which U.S. consumer gasoline expenditures are transferred abroad. It accounts for time variation in the share of gasoline imports in U.S. gasoline consumption as well as time variation in the share of U.S. net crude oil imports in the domestic use of crude oil. The underlying data have been seasonally adjusted.

Figure 2: Purchasing Power Shocks under Alternative Specifications, 1973.1-2016.3

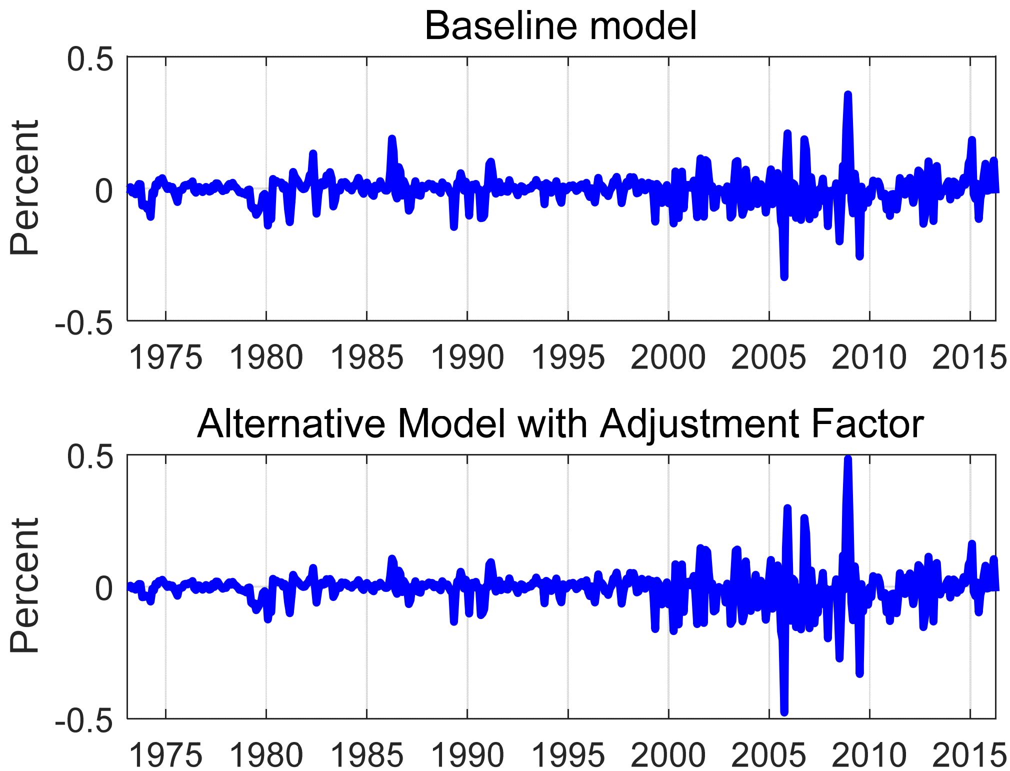

Notes: Authors' calculations based on BEA data. The adjustment factor is shown in Figure 1. The purchasing power shock used in the baseline model without loss of generality has been scaled by the average adjustment factor of 0.49, so the magnitude of the shocks can be compared directly.

CHRISTIANE BAUMEISTER
University of Notre Dame

LUTZ KILIAN
University of Michigan

Lower Oil Prices and the U.S. Economy: Is This Time Different?

ABSTRACT We explore the effect of the sharp and sustained decline after June 2014 in the global price of crude oil (and hence in the U.S. price of gasoline) on U.S. real GDP growth. Our analysis suggests that this decline produced a cumulative stimulus of about 0.9 percent of real GDP by raising private real consumption and non-oil-related business investment, and an additional stimulus of 0.04 percent, reflecting a shrinking petroleum trade deficit. This stimulative effect, however, has been largely offset by a large reduction in real investment by the oil sector. Hence, the net stimulus since June 2014 has been close to zero. We show that the U.S. economy's response was not fundamentally different from that observed after the oil price decline of 1986. Then as now, the U.S. economy's response is consistent with standard economic models of the transmission of oil price shocks. We find no evidence that frictions in reallocating capital and labor across sectors or increased uncertainty about the price of gasoline explain the sluggish response of U.S. real GDP growth. Nor do we find evidence of financial contagion, of spillovers from oil-related investment to non-oil-related investment, of an increase in household savings, or of households deleveraging.

Between June 2014 and March 2016, the real price of oil declined by 66 percent. There has been much debate about the effect of this sharp decline in global oil prices and, hence, in U.S. gasoline prices on U.S. growth. Many observers expected this oil price shock to boost the U.S. economy. Table 1 shows that, nevertheless, average U.S. real economic growth has increased only slightly since 2014:Q2, from 1.8 percent to 2.2 percent. Breaking down the components of real GDP reveals a striking discrepancy between sharply reduced average growth in real nonresidential

Table 1. Average Growth in U.S. Real GDP and Some of Its Components[a]

Component of real GDP	*2012:Q1–2014:Q2*	*2014:Q3–2016:Q1*
Real GDP	1.8	2.2
Private consumption	1.9	2.9
New motor vehicles	6.0	–2.8
Nonresidential investment	5.1	1.5
Oil-related investment[b]	7.2	–48.2
Non-oil-related investment	4.9	4.6
Exports	3.2	0.7
Petroleum exports	7.0	13.0
Imports	2.3	2.9
Petroleum imports	–7.7	3.8

Source: U.S. Bureau of Economic Analysis.
a. Growth is measured in percentage changes at annual rates.
b. Oil-related investment includes investment in petroleum and natural gas structures as well as mining and oil field machinery.

investment, driven by a dramatic decline in oil-related investment, and substantially higher average growth in real private consumption. Moreover, real petroleum imports, which had been falling before 2014:Q2, have been rising again since 2014:Q2, while the growth in real petroleum exports has nearly doubled, reducing the petroleum trade deficit and adding to real GDP growth. The increase in real petroleum exports is in contrast to the decline in overall real exports since 2014:Q2.

The evidence in table 1 raises a number of questions. Unexpected declines in the real price of oil may affect the U.S. economy, for example, to the extent that they lower firms' costs of producing domestic goods and services. Why, then, did the decline in the real price of oil not cause a strong economic expansion, as presumed in standard macroeconomic textbooks, which interpret oil price shocks as shifts of the domestic aggregate supply curve (or, in a more modern framework, as shifts in the cost of producing domestic real output)? Unexpected declines in the real price of oil also matter for the economy, because they increase the demand for other domestic goods and services, as consumers spend less of their income on motor fuels. One question of interest, thus, is by how much we would have expected private real consumption to increase as a result of the windfall income gain caused by lower oil prices. Did the actual growth in private real consumption match expected growth, or was it perhaps held back because the decline in the global price of crude oil was not fully passed on to retail fuel prices? Or did consumers simply choose not to spend their income gains, but to instead pay off their debts or increase their savings?

Finally, why did the real consumption of motor vehicles decline, despite an overall increase in private real consumption? Were consumers perhaps reluctant to buy new automobiles because of increased uncertainty about future gasoline prices, holding back overall economic growth?

Another puzzle is why growth in private nonresidential investment declined as much as it did after 2014:Q2. Clearly, the answer is related to the increased importance of U.S. shale oil production, raising the question of whether the growth of the shale oil sector has changed the transmission of oil price shocks to the U.S. economy. The decline in oil-related investment in response to falling oil prices not only has a direct effect on U.S. real GDP; there are also broader implications to consider. One concern has been that the decline of the shale oil sector may have slowed growth across oil-producing states, dragging down aggregate U.S. growth. Another conjecture has been that lower investment by oil producers may have slowed growth in other sectors of the economy nationwide, as the demand for structures and equipment used in oil production declined. A third conjecture has been that risky loans to oil companies may have undermined the stability of the banking system, disrupting financial intermediation. A related concern has been that the sustained decline in the real price of oil after 2014:Q2 may have caused an economic slowdown by leaving assets and oil workers stranded in a sector that is no longer competitive.

Equally surprising is the change in the petroleum trade balance since 2014:Q2, which does not conform to the conventional wisdom that an unexpected decline in the price of oil is associated with rising petroleum trade deficits, as domestic oil production declines. Finally, the substantial decline in U.S. real nonpetroleum exports is a reminder that the decline in the real price of oil itself was associated at least in part with a global economic slowdown that in turn needs to be taken into account in explaining the comparatively slow U.S. economic growth.

In this paper, we investigate the empirical support for each of these conjectures. We examine the channels by which the 2014–16 oil price decline might have affected the U.S. economy and assess their quantitative importance, drawing on a wide range of macroeconomic, financial, and survey data, at both the aggregate level and the sectoral and state levels. Our objective is to quantify how much of the evidence in table 1 can be explained by the unexpected decline in the real price of oil, without ruling out the possibility that other economic shocks may have affected U.S. economic growth at the same time. In section I, we provide evidence for the view that the demand channel of the transmission of oil price shocks to the U.S. economy is more important than the supply (or cost) channel emphasized in many theoretical models. This evidence motivates our emphasis on the demand channel of transmission throughout the remainder of this paper.

Our discussion of the demand channel focuses in particular on understanding the evolution of private consumption, investment spending, and the petroleum trade balance. In section II, we examine to what extent standard economic models of the transmission of oil price shocks that focus on changes in consumers' discretionary income, as the decline in oil prices is passed through to retail fuel prices, can explain the growth in real private consumption documented in table 1 (Edelstein and Kilian 2009; Hamilton 2009, 2013; Kilian 2014). 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. This reasoning is analogous to the conventional analysis of an unexpected increase in the real prices of oil and gasoline. In the words of Janet Yellen (2011):

> Higher oil prices lower American income overall because the United States is a major oil importer and hence much of the proceeds are transferred abroad.... Thus, an increase in the price of crude oil acts like a tax on U.S. households, and...tends to have a dampening effect on consumer spending....Staff analysis at the Federal Reserve Board indicates that a[n]...increase in retail gasoline prices...reduces household disposable income...and hence tends to exert a significant drag on consumer spending.

Yellen goes on to stress that the effect of these shocks on the economy has changed, as households' dependence on gasoline has evolved over time. Underlying this analysis is the view that oil price shocks represent terms-of-trade shocks that affect domestic spending and, hence, real GDP growth through a Keynesian multiplier. Although some of the so-called oil tax that is transferred abroad may ultimately be recycled, as oil-exporting countries directly or indirectly increase imports of goods and services produced in the United States, this petrodollar recycling tends to occur with a considerable delay, if at all.[1]

In response to an unexpected decline in the price of oil, as occurred after June 2014, the basic mechanism described by Yellen (2011) operates

1. As discussed by Hamilton (2013), an exogenous increase in the real price of oil may have real effects, even in a closed economy. Given that the price elasticity of gasoline demand is comparatively low, an exogenous increase in the price of gasoline causes a reduction in consumers' discretionary income. Although consumers' increased spending on gasoline represents income for someone else, by construction in a closed economy it may take considerable time for this income to be returned to consumers in the form of company profits, royalties, or dividends paid to shareholders, or to be spent by oil companies in the form of increased investment expenditures. Differences in the marginal propensity to spend thus may affect the overall level of spending and hence the business cycle in the short run.

in reverse, and is expected to generate a stimulus for the U.S. economy. We quantify this effect based on estimates of a linear regression model of the relationship between changes in real U.S. private consumption and changes in consumers' purchasing power associated with gasoline price fluctuations, controlling for the evolution of the share of fuel expenditures in total consumer expenditures. Estimates of this baseline model suggest that unexpectedly low oil prices cumulatively raised U.S. real GDP after 2014:Q2 by about 0.7 percent, as purchasing power increased and private consumption expanded. We show that similar estimates are also obtained after incorporating a measure of changes in the dependence of the U.S. economy on crude oil imports and gasoline imports in the construction of the purchasing power shocks.

In section III, we examine an alternative view in the literature, according to which the conventional linear model is overstating the stimulus for real GDP growth, because the true relationship between the price of oil and the economy is governed by a time-invariant, but nonlinear, process. Proponents of this view point to a number of indirect channels of transmission ignored by the baseline model. For example, it could be argued that the stimulative effects of the oil price decline discussed above are offset by delays in the reallocation of resources (Hamilton 1988; Bresnahan and Ramey 1993; Davis and Haltiwanger 2001; Ramey and Vine 2011; Herrera and Karaki 2015; Herrera, Karaki, and Rangaraju 2016) or by higher oil price uncertainty (Bernanke 1983; Pindyck 1991; Elder and Serletis 2010; Jo 2014). Either of these economic mechanisms would generate a nonlinearity that could explain why unexpected real oil price increases are recessionary, yet unexpected real oil price declines may not be followed by economic expansions and may even be recessionary. In section III, we provide both aggregate and disaggregate evidence suggesting that neither of these interpretations fits the recent episode.

In section IV, building on the results in sections II and III, we quantify the extent to which unexpectedly low oil and gasoline prices have stimulated private nonresidential investment (excluding the oil sector). We make the case that this stimulus can be estimated from a linear regression model similar to the model we utilized for private consumption. This investment stimulus adds another 0.2 percent in cumulative real GDP growth to the consumption stimulus of 0.7 percent.

A common view is that the relationship between the economy and changes in the price of oil has changed in recent years, calling into question estimates of the stimulus based on linear regression models. Proponents of this view would argue that this latest episode of declining oil prices is fundamentally different from previous episodes of sustained declines in the price of oil, such as the 1986 episode, so nothing about the economy's response can be learned from fitting regressions to historical data. One candidate explanation for such a structural shift in recent years is the increased importance of the U.S. shale oil sector since late 2008, which created potentially important additional effects of oil price shocks on domestic value added, aggregate nonresidential investment expenditures, the petroleum trade balance, and the stability of the banking sector. Likewise, a structural shift could arise if consumers used the windfall income associated with lower oil prices to reduce their mortgage debt and credit card debt rather than spending the extra income, as in years past. In section V, we examine the empirical evidence for these and other hypotheses. We find no evidence that households' savings behavior has changed or that households have been deleveraging, but we find evidence of an unprecedented decline in oil-related investment in the U.S. economy, and of a systematic reduction in net petroleum imports. The latter two structural shifts complicate the task of assessing the U.S. economy's response to the recent decline in the price of oil.

A simple national income accounting calculation in section VI suggests that the stimulative effect of lower oil prices on private real consumption, non-oil-related nonresidential investment, and net petroleum exports after June 2014 was approximately offset by the reduction in real investment by the U.S. oil sector. The net stimulus raised average real GDP growth by a paltry 0.2 percentage point at annual rates. Finally, in section VII, we compare the economy's response to the decline in the oil price after June 2014 with its response to the 1986 oil price decline, and make the case that there are more similarities than differences. The most important difference is that the recent decline in the real price of oil was about twice as large as the decline in 1986, causing a sharper contraction in oil investment than in 1986. Moreover, unlike the 1986 oil price decline, it was associated in part with a global economic slowdown, reflected in a substantial decline in the growth of U.S. real nonpetroleum exports, without which average U.S. real GDP growth is likely to have reached 2.5 percent at annual rates after 2014:Q2.

I. How Important Is the Cost Channel of the Transmission of Oil Price Shocks?

The traditional undergraduate textbook analysis of the effects of oil price shocks on oil-importing countries equates lower oil prices with a reduction in the cost of producing domestic goods and services (and, hence, with a

Figure 1. Traffic Volumes in the U.S. Transportation Sector, January 2013–March 2016[a]

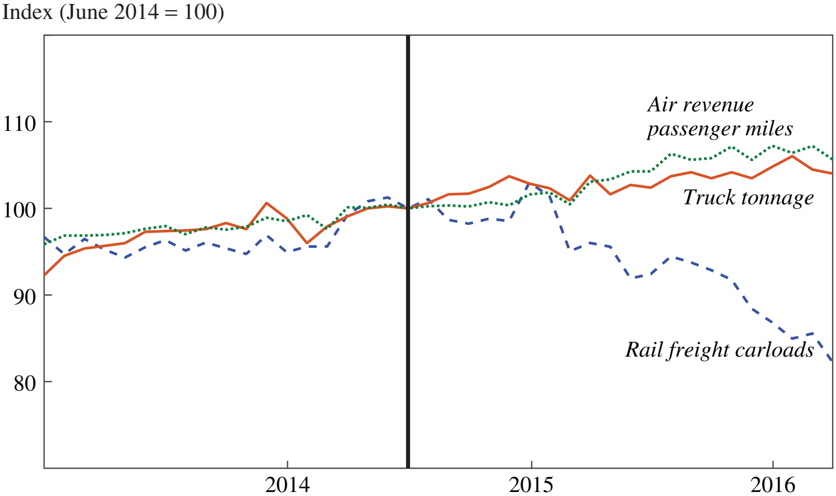

Source: U.S. Bureau of Transportation Statistics.
a. The vertical line marks June 2014, which is the month before the oil price decline began.

shift in the domestic aggregate supply curve along the domestic aggregate demand curve). This view has merit, to the extent that firms directly or indirectly rely on oil (or oil-based products) as a major factor of production. Examples of such industries include the transportation sector, some chemical companies, and rubber and plastics producers. For most industries, however, this channel is not likely to be important. In fact, a large share of the oil used by the U.S. economy is consumed by final consumers rather than by firms, which explains why more recent studies have typically interpreted oil price shocks as affecting the disposable income of consumers. This more contemporary view implies that oil price shocks are primarily spending or demand shocks for the U.S. economy. Within the traditional undergraduate textbook model, they can be thought of as shifts in the aggregate demand curve along the aggregate supply curve.

Some informal evidence regarding the relative importance of the supply (or cost) channel of the transmission of oil price shocks and the demand channel of transmission may be obtained by examining which sectors benefited and which suffered after the oil price decline. For example, there is general agreement that transportation is the sector most sensitive to changes in fuel prices. The data provide at best partial support for this view. Figure 1 shows that the volume of truck tonnage evolved largely along the same trend line before and after June 2014. In contrast, airline passenger traffic accelerated, but only with a delay of half a year, which likely reflects the fact that airlines had hedged against higher oil prices in futures markets and were able to pass on these added costs to the retail customer when the price of oil fell. Rail freight traffic initially remained relatively stable, but fell starting in early 2015, reflecting the global economic slowdown in general, and a substantial decline in U.S. coal shipments in particular. To a lesser extent, this pattern is also found in barge traffic and air freight traffic. A much smaller decline in rail passenger traffic, in contrast, is likely to reflect substitution away from trains and toward automobiles. Overall, these effects appear modest at best, and they are at odds with the view that lower fuel costs have a large effect on real output in the transportation sector.

This conclusion is corroborated by data on the excess stock returns for selected sectors and individual firms relative to the overall U.S. stock market index between July 2014 and March 2016.[2] All results are expressed as average excess returns at annual rates. In general, companies that cater to U.S. consumers tend to appreciate in value more than the average company. In particular, candy and soda (+7 percent), beer and liquor (+10 percent), and tobacco (+16 percent) do well, perhaps because such goods are sold at gas stations; but food products (+7 percent), and apparel (+11 percent) also do well. Both tourism (+11 percent) and restaurants, hotels, and motels (+8 percent) benefited from lower oil prices, as consumer demand rose. So did retail sales (+14 percent). Amazon (+38 percent) and The Home Depot (+32 percent) did particularly well. Only recreation, entertainment services, and publishing did not partake in this boom.

Unsurprisingly, the petroleum and natural gas sector (–28 percent) was hit hard. Within this sector, refining companies that use crude oil as a production input fared somewhat better. Other industries that rely on oil as a major input and hence would have been expected to profit from lower oil prices—such as rubber and plastics (+4 percent) and logistics (+2 percent)—did not benefit much, and chemicals (–6 percent) actually performed worse than the overall market, arguing against an important supply (or cost) channel of transmission. Airlines (+15 percent) benefited both from lower fuel costs and higher travel demand. Likewise, textiles

2. The analysis is based on individual returns from Bloomberg and value-weighted, industry-level stock returns obtained from http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. The benchmark portfolio is the value-weighted S&P 500 stock return from Bloomberg.

were helped by lower input costs and higher demand (+13 percent). The surprising fact that auto companies performed below average (–9 percent) is largely explained by weak foreign sales, reflecting the recent global economic slowdown. Sectors tied to commodity markets, such as agriculture (−12 percent) and mining (−31 percent), performed poorly for the same reason. Steel (−26 percent), fabricated metal products (−51 percent), machinery (−19 percent), and shipbuilding and railroad equipment (−13 percent) all suffered from lower demand, mainly due to the decrease in global real economic activity.

We conclude that the supply (or cost) channel, which is emphasized in many theoretical models of the transmission of oil price shocks developed in the 1980s and 1990s, may be safely neglected. Lower fuel costs do not appear to provide much of a stimulus to firms that are oil-intensive in production. The few sectors other than refining that are heavily dependent on oil inputs performed only marginally better than the rest of the economy after June 2014, if at all. In contrast, sectors sensitive to fluctuations in consumer demand did far better than average, lending support to the conventional view among policymakers and oil economists that the demand channel of the transmission of oil price shocks to the U.S. economy is more important than the supply channel (Kilian 2014). Our industry-level analysis of excess stock returns provides strong evidence of a stimulus to U.S. consumer demand, but also of lower demand stemming from a global economic slowdown, which corroborates related results in the literature, including the narrative evidence given by Lee and Ni (2002) and the regression evidence given by Kilian and Park (2009).

II. How Much Did the Unexpected Decline in the Price of Oil Stimulate Consumption?

Given the evidence presented in section I, our analysis focuses on the demand channel of transmission. We first examine private consumer spending, which accounted for 69 percent of U.S. GDP in 2014. For the oil price decline after 2014:Q2 to have stimulated U.S. private consumption, it was necessary for this decline to have been passed through to retail fuel prices. We therefore first quantify the extent to which U.S. gasoline prices have declined in response to lower crude oil prices, taking account of the cost share of crude oil in producing gasoline. The answer to this question is not obvious because there is a long-standing view that oil price declines are not necessarily passed on to retail gasoline prices as quickly as oil price increases (Venditti 2013). We provide evidence that the pass-through is

Figure 2. U.S. Price of Gasoline and Cost of Crude Oil, January 2000–March 2016

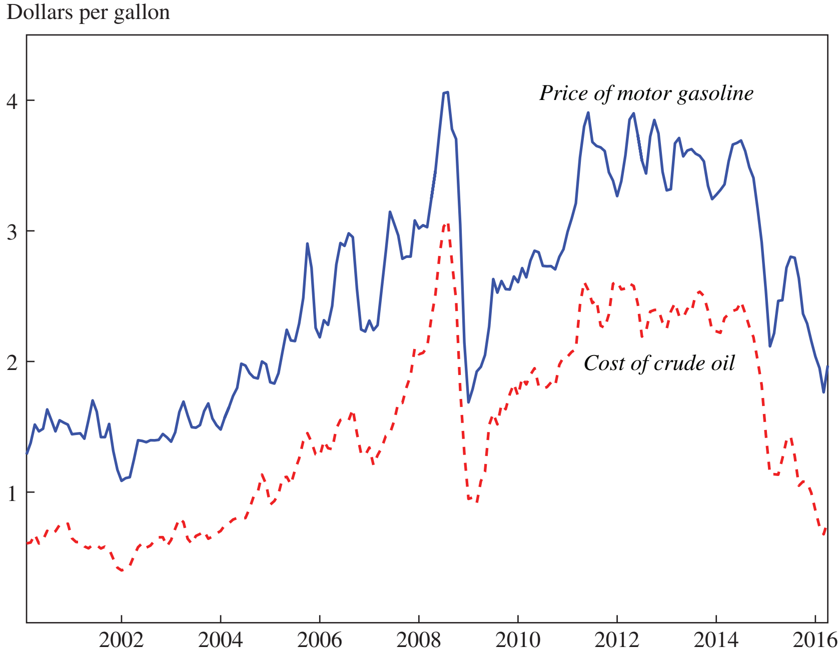

Source: U.S. Energy Information Administration, *Gasoline and Diesel Fuel Update*, *Gasoline Pump Components History*.

symmetric and that the recent oil price decline has been fully passed on to retail gasoline prices. We then quantify the changes in U.S. consumers' purchasing power associated with unexpected changes in the price of gasoline and estimate the cumulative effect of these shocks on real private consumption, controlling for changes in the share of gasoline expenditures in total consumer expenditures. The magnitude of the estimated stimulus is shown to be consistent with a back-of-the-envelope calculation that treats the change in the gasoline price as taking place, all else equal, and takes account of the price elasticity of gasoline demand.

II.A. Has the Decline in the Price of Oil since June 2014 Been Passed Through to Gasoline Prices?

Figure 2 shows the price of gasoline at the pump and the cost of the crude oil used in producing gasoline. The difference between these time series reflects changes in gasoline taxes and in the costs of refining crude oil and

Figure 3. Index of U.S. Price of Gasoline and Cost of Crude Oil, June 2014–March 2016

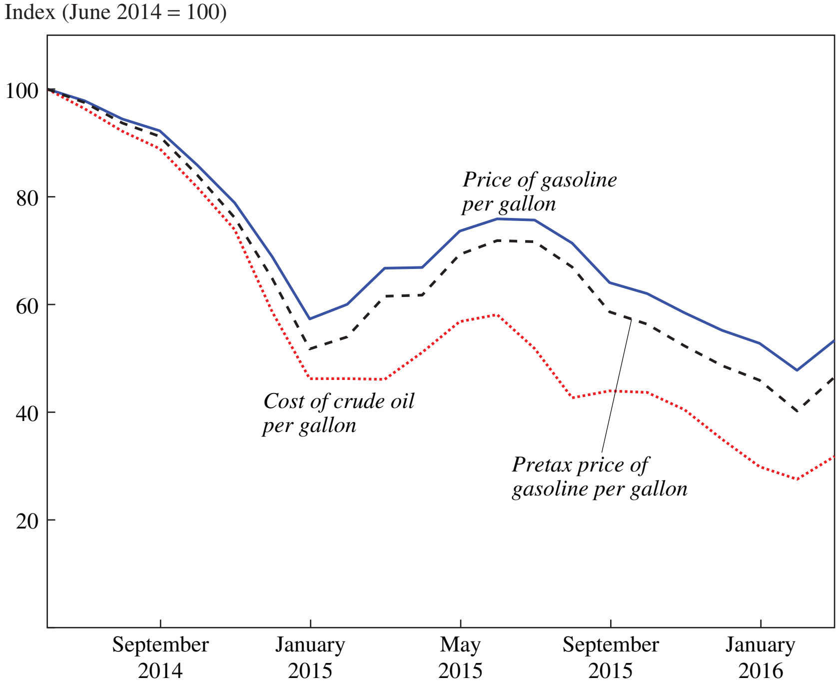

Source: U.S. Energy Information Administration, *Gasoline and Diesel Fuel Update*, *Gasoline Pump Components History*.

of marketing and distributing gasoline.[3] Figure 3 zooms in on events since June 2014. All data are expressed as index numbers normalized to equal 100 in June 2014. Between June 2014 and December 2015, the price of gasoline fell by 45 percent, whereas the cost of crude oil fell by 65 percent (about as much as the spot price of Brent crude oil). Some of that difference is accounted for by a slight increase in gasoline taxes, but even the pretax

3. Although there is a large degree of comovement between the cost of oil and the price of gasoline, this comovement is by no means perfect. For example, in 2005, when Gulf Coast oil refiners were forced to shut down due to Hurricane Rita and Hurricane Katrina, causing a refining shortage, there was a sharp spike in the price of gasoline, but not in the price of crude oil, illustrating that occasionally changes in gasoline prices are not just determined by changes in the price of crude oil (Kilian 2010). A regression of the price of retail gasoline on an intercept and the cost of crude oil, as shown in figure 2, yields a slope coefficient of 1.1, suggesting a nearly one-for-one relationship in the long run.

Table 2. Evidence of Pass-Through from Oil Price to Gasoline Price by Episode

Percent	*January 2007–July 2008*	*July 2008–December 2008*	*December 2008–April 2011*	*June 2014–March 2016*
Change in U.S. retail gasoline price	81.3	–58.5	125.3	–46.7
Change in the cost of crude oil used in producing a gallon of U.S. gasoline	155.0	–69.2	175.4	–68.2
Change in the Brent price of crude oil	147.2	–69.9	208.5	–65.8
Average cost share of crude oil in U.S. gasoline production	63.3	65.2	64.6	51.4
Expected change in U.S. gasoline price[a]	98.1	–45.1	113.3	–35.0

Source: U.S. Energy Information Administration, *Gasoline and Diesel Fuel Update*, *Gasoline Pump Components History*.

a. The expected percent change in the U.S. gasoline price is constructed by weighting the percent change in the dollar cost of crude oil used in producing a gallon of gasoline by the average cost share of crude oil.

price of gasoline only fell by 53 percent. At first sight, this evidence might seem to imply that refiners or gasoline distributors failed to pass on the full cost savings resulting from the 2014–16 oil price decline to consumers. It is important to keep in mind, however, that historically only about half the price of gasoline has consisted of the cost of crude oil, so even with perfect pass-through one would expect a percentage decline in the price of gasoline only about half as large as the percentage decline in the cost of crude oil.

Table 2 examines the extent to which cumulative changes in the cost of the crude oil used in producing gasoline have been reflected in changes in the price of gasoline based on four key episodes, two of which involve increases in the cost of crude oil and two of which involve declines. For example, between January 2007 and July 2008, the cost of crude oil increased cumulatively by 155 percent (slightly more than the spot price of Brent crude oil). Given an average cost share of 63.3 percent over this period, all else equal, one would have expected the price of gasoline to cumulatively increase by 98.1 percent. The actual cumulative increase, 81.3 percent, was somewhat lower, but not far from this benchmark. Another large cumulative increase in the cost of oil occurred between December 2008 and April 2011. The cost of oil surged by 175.4 percent (somewhat less than the spot price of Brent crude oil). Given the average cost share of 64.6 percent, one would have expected the price of gasoline, all else equal, to increase by 113.3 percent. The actual increase by 125.3 percent was somewhat higher, but in the same ballpark.

What about declines in the cost of crude oil? Between July 2008 and December 2008, the cost of crude oil fell by 69.2 percent cumulatively, which, given the average cost share of 65.2 percent, would have led us to expect the gasoline price to decline by 45.1 percent, somewhat less than the observed decline of 58.5 percent. Likewise, the cumulative decline in the cost of oil of 68.2 percent between June 2014 and March 2016, given the average cost share of 51.4 percent, translates to an expected decline of 35 percent in the U.S. gasoline price, compared with a somewhat larger decline of 46.7 percent in the data.

These four examples are consistent with the view that, on average, the observed changes in gasoline prices are roughly as large as one would have expected under the assumption of perfect pass-through, given that gasoline prices may vary for a range of other reasons ranging from refinery outages to changes in the retail market structure (Baumeister, Kilian, and Lee 2016). The decline in the price of gasoline that occurred in 2014–15, if anything, exceeded what one would have expected based on the pass-through from the cost of oil to the gasoline price at the pump. There is no evidence of asymmetries in the pass-through between declines and increases in the cost of oil in table 2, which corroborates the econometric results obtained by Venditti (2013).

II.B. How Has the Consumption of Gasoline Evolved since June 2014?

Lower gasoline prices increase the discretionary income of consumers to the extent that the same amount of gasoline may be purchased with less income. Lower gasoline prices, however, also provide an incentive to increase gasoline consumption that reduces the extra income available for other purchases. Figure 4 shows the evolution of seasonally adjusted U.S. gasoline consumption—defined as the sum of the motor gasoline consumed by the industrial, commercial, and transportation sectors—since June 2014. Gasoline consumption cumulatively increased by 5.5 percent between June 2014 and January 2015, reaching 7.4 percent by March 2016. The increase in gasoline consumption coincided with a 5 percent increase in vehicle miles

Figure 4. Cumulative Change in U.S. Gasoline Consumption, June 2014–March 2016[a]

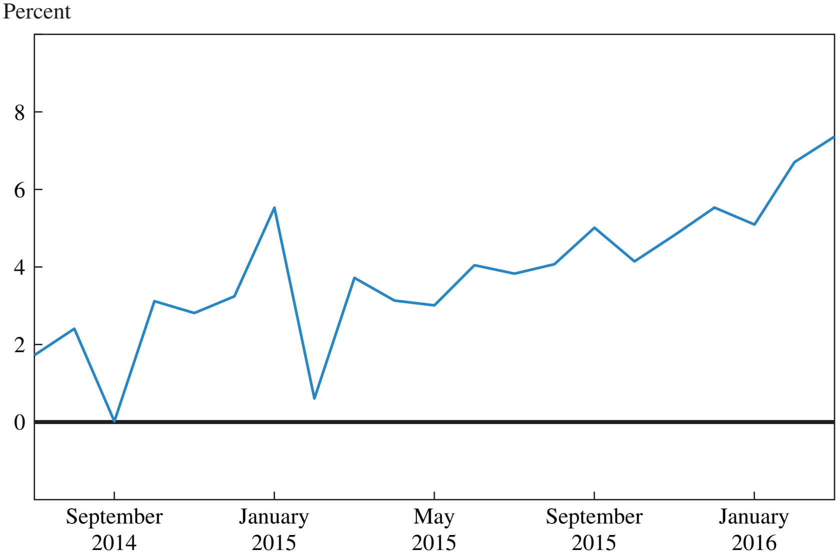

Source: U.S. Bureau of Economic Analysis; U.S. Bureau of Labor Statistics.

a. Data are based on the sum of transportation, industrial, and commercial gasoline consumption, seasonally adjusted by the authors.

traveled since June 2014, as shown in figure 5. At the same time, the fuel economy of new cars and light trucks, as measured by the average sales-weighted miles per gallon reported by the University of Michigan's Transportation Research Institute, fell by 2 percent, from a peak of 25.8 miles per gallon in August 2014 to 25.3 miles per gallon in March 2016, reflecting changes in the composition of new vehicles.

II.C. Measuring Gasoline Price Shocks

Gasoline price shocks are defined as the difference between what the price of gasoline was expected to be ex ante and what it actually turned out to be. In recent work, we have made the case that what matters when quantifying gasoline price shocks is the expectation of the decisionmaker whose behavior one seeks to understand (Baumeister and Kilian 2016a). If we want to understand the response of U.S. consumers, for example, the relevant measure of gasoline price expectations is consumers' own expectations, no matter how inaccurate this measure may be by statistical criteria. The Michigan Survey of Consumers provides data starting in February 2006 for consumers' expectations of the change in gasoline prices over

Figure 5. Moving 12-Month Vehicle Miles Traveled, January 2009–March 2016[a]

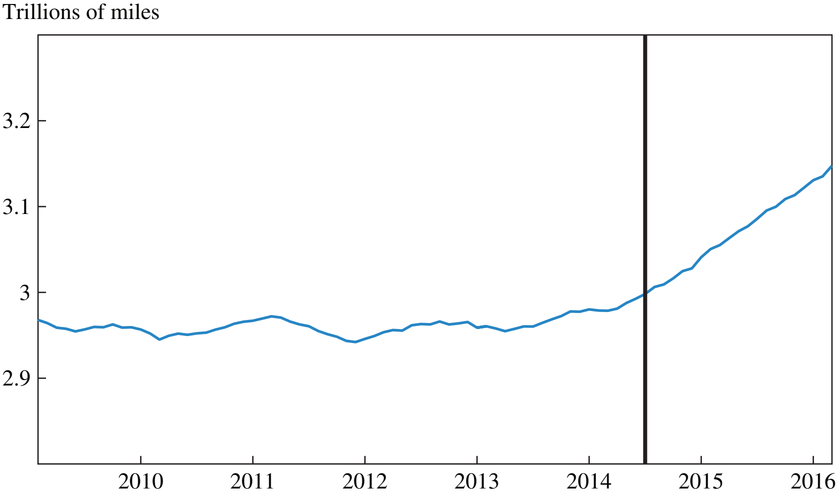

Source: U.S. Bureau of Transportation Statistics, Federal Highway Administration.
a. The vertical line marks June 2014, which is the month before the oil price decline began.

the next 12 months. Based on these data, Anderson, Kellogg, and Sallee (2013) document that consumers, with rare exceptions, expect the nominal price of gasoline to grow at the expected rate of inflation. An obvious question is whether this approximation remains valid even during a decline in the price of gasoline as sustained as the decline that started in June 2014.

We address this question in figure 6, which plots the expectation of the price of gasoline implied by the survey data. The gasoline price expectation is constructed by adding the median expected change in gasoline prices over the next 12 months from the Michigan Survey of Consumers to the average U.S. price of gasoline from the U.S. Energy Information Administration's (EIA's) *Monthly Energy Review.* Figure 6 shows that this survey measure closely tracks the no-change forecast of the real price of gasoline, adjusted for the median expected change in the price level over the next 12 months, as reported in the Michigan Survey of Consumers, even after June 2014. This evidence suggests that one can approximate consumers' expectations of the real gasoline price based on a simple no-change forecast of the real price of gasoline. We employ this approach to construct a monthly time series of the real gasoline price shocks experienced by U.S. consumers from January 1970 to March 2016.

Figure 6. Gasoline Price Expectations, One Year Ahead, February 2006–March 2016[a]

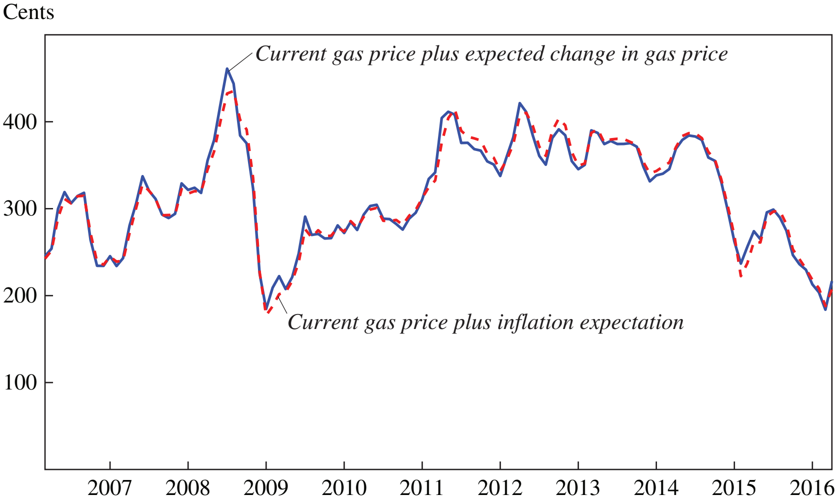

Source: Michigan Survey of Consumers; U.S. Energy Information Administration, *Monthly Energy Review*.

a. The gasoline price expectation is obtained by adding the median expected change in gasoline prices over the next 12 months from the Michigan Survey of Consumers to the average U.S. price of gasoline from the *Monthly Energy Review*. The survey measure closely tracks the no-change forecast of the price of gasoline adjusted for the median expected change in the price level over the next 12 months in the Michigan Survey of Consumers, as previously noted by Anderson, Kellogg, and Sallee (2013).

Let $S_t \equiv (R_t^{gas} - E_{t-1} R_t^{gas})/E_{t-1} R_t^{gas}$, where R_t^{gas} is the real price of gasoline, defined as the average nominal price of gasoline and other motor fuel, P_t^{gas}, as reported by the U.S. Bureau of Economic Analysis (BEA), deflated by the overall personal consumption expenditures deflator, P_t^C, and $E_{t-1} R_t^{gas} = R_{t-1}^{gas}$. This shock measure simply corresponds to the percentage change in the real price of gasoline and other motor fuel, as shown in the upper panel of figure 7. How much this gasoline price shock matters to U.S. consumers depends on the share of expenditures on gasoline and other motor fuels in overall consumer expenditures. For a given unexpected increase in the real price of gasoline, the higher this expenditure share, the higher the potential reduction in consumers' discretionary income, because income spent on gasoline cannot be spent on other goods and services. As illustrated in the middle panel of figure 7, this share has fluctuated between about 2 and 5 percent since 1970. In mid-1973, in early 2006, and again in mid-2014, this share was near its long-run average value of 3 percent.

A measure of the shock to consumers' purchasing power may be constructed as

$$PP_t \equiv -S_t \times \frac{C_t^{gas} P_t^{gas}}{C_t P_t^C},$$

where C_t^{gas} is real U.S. gasoline consumption and C_t is real total consumption, as reported by the BEA. The series of purchasing power shocks, PP_t, is shown in the bottom panel of figure 7. It is the latter shock series to which consumers respond, rather than the gasoline price shock in the upper panel. Figure 7 shows clear evidence of an unexpected increase in purchasing power in 1986, following a sharp drop in the global price of crude oil; it shows repeated, unexpected reductions in purchasing power between 1999 and 2008 during the surge in global oil prices; a large positive purchasing power shock in late 2008, associated with the financial crisis, that was quickly reversed in early 2009; and a series of positive and negative purchasing power shocks since June 2014, during the period of interest in this paper.

II.D. The Baseline Linear Model

The question of ultimate interest is by how much these purchasing power shocks stimulated real private consumption. Our analysis is based on a monthly model that embodies the identifying assumption that changes in purchasing power are predetermined with respect to real consumption.[4] Let Δc_t denote the percentage change in monthly real consumption (demeaned to account for the drop in average consumption growth from 3.3 percent at annual rates to 2.1 percent after December 2008), and let PP_t denote the monthly shock to consumers' purchasing power, as defined in subsection II.C. The shocks are normalized such that a positive shock indicates an increase in purchasing power. Then the response of consumption to purchasing power shocks may be estimated from the ordinary least squares regression

$$\Delta c_t = \sum_{i=1}^{6} \beta_i \Delta c_{t-i} + \sum_{i=0}^{6} \gamma_i PP_{t-i} + u_t, \tag{1}$$

4. For related approaches see, for example, Edelstein and Kilian (2009) and Hamilton (2009). The validity of this identifying assumption is supported by evidence in Kilian and Vega (2011).

Figure 7. Measuring Shocks to Consumers' Purchasing Power, February 1970–March 2016

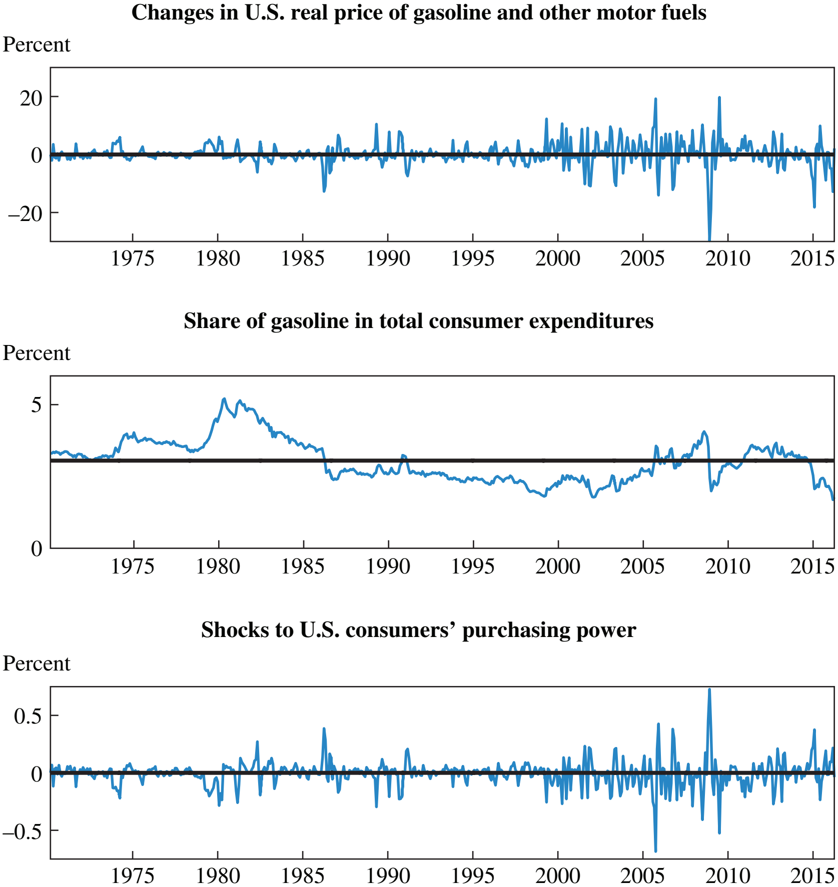

Sources: U.S. Bureau of Economic Analysis; authors' calculations.

where u_t denotes the regression error.[5] Given that there has been considerable variation in the magnitude and sign of the changes in purchasing power since June 2014, a more useful approach to studying the changes in U.S. real private consumption over this period is to compute the cumulative effect of these purchasing power gains and losses on real consumption

5. The lag order choice follows Edelstein and Kilian (2009). The estimation sample is February 1970 to March 2016.

Table 3. Predicted Cumulative Percent Change in U.S. Real Consumption, July 2014–March 2016[a]

Component	*Percent*
Total consumption	1.2
Durables	2.9
New motor vehicles	6.7
Nondurables	0.8
Services	0.8

Source: Authors' calculations.

a. Historical decomposition based on the fitted values of the baseline regression model in equation 1. The estimation sample is February 1970 to March 2016.

since June 2014. Table 3 shows that, according to the model, purchasing power shocks cumulatively stimulated U.S. real private consumption by 1.2 percent and account for most of the observed 1.3 percent cumulative increase in total real private consumption, relative to trend, between July 2014 and March 2016. Taking account of the drift, the model predicts an average growth rate of 2.8 percent in real private consumption at annual rates, compared with 2.9 percent in table 1.

Part of the estimated cumulative increase in consumption is accounted for by the operating cost effect, which refers to an increase in purchases of automobiles in response to unexpectedly lower gasoline prices. This operating cost effect amplifies the overall consumption response over and above the discretionary income effect (Hamilton 1988). Table 3 confirms the existence of a disproportionately larger stimulus of nearly 3 percent for durables (which in turn is largely driven by the consumption of new motor vehicles). Weighting the 6.7 percent stimulus for the consumption of new motor vehicles in table 3 by the share of new motor vehicles in private consumption of 2.3 percent suggests a cumulative operating cost effect of 0.15 percent. Given the overall cumulative consumption response of 1.2 percent in the baseline model, this implies a discretionary income effect of about 1.05 percent.[6]

6. Ramey and Vine (2011) propose scaling the nominal gasoline price during December 1973–May 1974 and during May 1979–July 1979 by a multiplicative factor intended to capture the waiting cost at gas stations associated with government-imposed gasoline price ceilings. Because the waiting cost is not associated with a transfer of income abroad, this adjustment must not be used in quantifying the discretionary income effect. It may affect the operating cost effect, however. Further sensitivity analysis shows that adjusting PP_t for the waiting cost only affects the third decimal place of our estimates of the operating cost, so the waiting-cost adjustment may be safely ignored.

A simple back-of-the-envelope calculation suggests that the magnitude of this estimate of the discretionary income effect is reasonable. The real price of gasoline and other motor fuels declined by 44.94 percent between June 2014 and March 2016. The share of gasoline expenditures in total expenditures in June 2014 was 3.17 percent. This allows consumers to purchase the same goods for a fraction of their income and frees up 1.13 percent of consumers' income for additional purchases:

$$(1-0.0317)\times 1+0.0317\times(1-0.4494)(1+0.37\times 0.4494)=0.9887,$$

where we estimated the increase in gasoline consumption based on the estimate of the price elasticity of gasoline demand of −0.37 reported by Coglianese and others (2016). This exercise suggests a discretionary income effect on consumption close to the estimate of 1.05 percent implied by the baseline model.

II.E. An Alternative Linear Specification

The regression model in equation 1 is designed to capture the extent to which discretionary income is injected into the U.S. economy or removed from the U.S. economy, as the terms of trade vary in response to oil price shocks. This model implicitly assumes that the share of the proceeds from gasoline that goes abroad is the same over time. To the extent that this share varies over time, the model provides only an approximation. It may seem that variation in the dependence of the U.S. economy on oil and gasoline imports over time would render this approximation inaccurate.

Assessing the empirical content of this concern is not straightforward. For example, it may be tempting to answer this question by testing for structural breaks in the parameters of model 1, but that approach would not be informative. It is well documented that mechanical applications of tests for structural stability on subsamples are prone to generating spurious rejections of the null hypothesis of a stable relationship between macroeconomic aggregates and oil or gasoline prices (for example, Kilian 2009; Kilian and Park 2009). Although the average responses of real consumption to purchasing power shocks may be reliably estimated using long samples, when considering short subsamples, these responses will change in magnitude and even in sign, as the composition of oil demand and oil supply shocks changes over time, giving the mistaken appearance of structural instability, even when there is no structural change at all. Spurious evidence of structural breaks arises whenever oil price fluctuations over a subsample are not representative of the full sample.

An alternative and more direct way to quantify the importance of changes in the dependence of the U.S. economy on oil and gasoline imports is to incorporate these changes in the construction of PP_t. A simple approximation is to weight U.S. consumer expenditures on gasoline by the share of the proceeds going abroad, resulting in an alternative definition of purchasing power shocks,

$$PP_t^{alternative} \equiv -S_t \times \frac{C_t^{gas} P_t^{gas}}{C_t P_t^C} \left(s_t^{gas\ imports} + \left(1 - s_t^{gas\ imports}\right) s_t^{net\ oil\ imports} \right),$$

where $s_t^{gas\ imports}$ is the seasonally adjusted share of U.S. motor gasoline imports in total U.S. motor gasoline consumption, as reported by the EIA, and $s_t^{net\ oil\ imports}$ is the seasonally adjusted share of U.S. net crude oil imports in the total use of crude oil by the U.S. economy, as defined by Kilian (2016a).[7] These data are available starting in January 1973. When estimating this alternative model, the implied overall consumption stimulus is 0.92 percent, which is somewhat lower than the 1.2 percent in the baseline specification but still in the same ballpark, adding credence to the baseline specification. Moreover, the operating cost effect is 0.14 percent in the alternative model, compared with 0.15 percent in the baseline model.

It can be shown that all substantive results in this paper are unaffected by the choice between the baseline model and the alternative model. We therefore focus on the baseline model in the remainder of the paper. The cumulative increase in real GDP growth implied by the combined effect of higher discretionary income and lower operating costs in the baseline model is 0.7 percent over the course of seven quarters, given the share of consumption in GDP of 69 percent and assuming a marginal import propensity of 15 percent. This conclusion is also consistent with a marked improvement in consumers' long-term expected business conditions, following the decline in the real price of oil. In the next sections, we examine the evidence for nonlinearities and structural breaks in the transmission of the oil price shocks as well as other factors that are not captured by this baseline model.

7. This measure is not only more relevant for understanding the foreign cost share of U.S. gasoline than the share of net imports in products supplied reported by the EIA, but it also avoids the ad hoc aggregation of crude oil and refined products. It is nevertheless only an approximation because it ignores changes in oil and gasoline inventories, assumes that the net share of imported crude oil is the same in the production of all refined products, does not differentiate between gasoline and other motor fuel, and makes no allowance for changes over time in the extent of petrodollar recycling from abroad.

III. Does the U.S. Economy Respond Asymmetrically to Unexpected Oil Price Increases and Decreases?

There is general agreement among economists on the existence of a discretionary income effect, but some economists have suggested that the effects of unexpectedly low oil prices are likely to be negligible, because the stimulative effects are offset by costly reallocations of resources or by higher uncertainty about gasoline prices. This view implies that the economy responds asymmetrically to unexpected increases and decreases in gasoline prices. The rationale for asymmetric responses of real output to oil price shocks hinges on the existence of additional indirect effects of unexpected changes in the real price of oil. There are two economic models that generate such indirect effects. One is the reallocation model of Hamilton (1988), which is the focus of subsection III.A; the other is the real options model of Bernanke (1983), which is discussed in subsection III.B. Next, we examine whether these models provide a plausible explanation for the sluggish growth of the U.S. economy following the decline in the price of oil after June 2014.

III.A. Did Frictions in Reallocating Capital and Labor Offset the Stimulus?

Relative price shocks, such as shocks to the real price of gasoline, can be viewed as allocative disturbances that cause sectoral shifts throughout the economy. For example, increased expenditures on energy-intensive durables such as automobiles in response to unexpectedly low real gasoline prices tend to cause a reallocation of capital and labor toward the automobile sector. Because the dollar value of such purchases may be large relative to the value of the fuel they use, even small changes in the relative price of gasoline can have potentially large effects on demand. This operating cost effect was discussed in section II. A similar reallocation may occur within the automobile sector, as consumers switch toward less fuel-efficient vehicles (Bresnahan and Ramey 1993). If capital and labor are sector specific or product specific and cannot be moved easily to new uses, these intersectoral and intrasectoral reallocations will cause labor and capital to be idle, resulting in cutbacks in real output and employment that go beyond the direct effects of a real gasoline price shock. For example, workers may be ill equipped to take different jobs without extensive job retraining. The same effect may arise if unemployed workers simply choose to wait for conditions in their sector to improve.

This indirect effect tends to amplify the direct recessionary effect on real output and unemployment of unexpected increases in the real price of gasoline, while dampening the economic expansion caused by unexpected declines in the real price of gasoline. There is a large empirical literature on potential asymmetries in the economy's response to positive and negative oil price shocks (for example, Herrera, Lagalo, and Wada 2011, 2015; Herrera and Karaki 2015; Kilian and Vigfusson 2016). Although the evidence thus far has not been supportive of models implying strongly asymmetric responses at the aggregate level, there have been comparatively few episodes of large oil price declines, so this latest episode provides an opportunity to take a fresh look at the evidence.

Given the challenges of measuring movements of capital across sectors, our discussion focuses on the movements of labor. Even in the latter case, it is difficult to directly assess the evidence for frictions. This would involve tracking workers after they lose their jobs in one sector. Some insights, however, may be gleaned from U.S. unemployment data at the aggregate level. If the hypothesis of frictional unemployment were empirically relevant, one would expect aggregate unemployment to increase relative to the level that would have prevailed in the absence of the decline in the price of gasoline. Such an effect would presumably manifest itself in an increase in the unemployment rate or, at the very least, a noticeably slower decline in the unemployment rate. Figure 8 shows that both the U.S. unemployment rate and the median duration of unemployment have been dropping steadily since late 2011. 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.

This pattern is by no means unprecedented. For example, figure 8 shows that the large and sustained decline in the price of gasoline after December 1985 was followed by a decline in the unemployment rate of a magnitude similar to the decline in the unemployment rate after June 2014. Table 4 compares the cumulative decline in the unemployment rate and in the median duration of unemployment that took place during these two episodes. Although the cumulative change in the real price of gasoline in the more recent episode was larger, the 0.96 percent cumulative gain in purchasing power over the first seven months was only slightly larger than the 0.85 percent increase observed in 1986, and so was the cumulative decline in the unemployment statistics. Then as now, there is no evidence of an increase in unemployment relative to trend. This evidence

Figure 8. U.S. Unemployment Data, January 1980–March 2016[a]

Source: U.S. Bureau of Labor Statistics.
a. The vertical lines mark December 1985 and June 2014, the months before the oil price declines began.

Table 4. Cumulative Changes in U.S. Unemployment Statistics Following the 1986 and 2014 Oil Price Declines

	January 1986–September 1987		*July 2014–March 2016*	
Statistic	*Absolute change*	*Relative change (percent)*	*Absolute change*	*Relative change (percent)*
Real gasoline price		−20.8		−44.9
Unemployment rate	−1.1 percentage points	−15.7	−1.1 percentage points	−18.0
Median unemployment duration	−0.8 weeks	−11.8	−2.1 weeks	−15.6

Sources: U.S. Bureau of Labor Statistics; U.S. Energy Information Administration.

casts some doubt on the view that the comparatively slow U.S. real GDP growth since June 2014 reflected frictions in the reallocation of labor.[8]

Further insights may be gained from employment data for the oil industry and related industries. Between December 2009 and its peak in October 2014, employment in this sector (defined as oil and natural gas extraction, including support activities and the construction of mining and oil field machinery and pipelines) increased by 278,000 workers. Between October 2014 and March 2016, employment fell by 166,000 workers. At the national level, the reduction in employment in the oil sector, although large in percentage terms, is clearly too small to matter much for the unemployment rate. Nor did the 2014 oil price decline have a large effect on net employment changes (Herrera, Karaki, and Rangaraju 2016).

We can get a better sense of how quickly these unemployed workers were absorbed by focusing on selected oil-producing states, such as Texas and North Dakota. For example, it has been suggested that the 1986 recession in Texas was caused by frictions impeding the reallocation of labor from the oil sector to other sectors. If so, one would also expect a pronounced increase in unemployment in Texas after June 2014. As of June 2014, the mining and logging sector accounted for 2.7 percent of nonfarm employment in Texas. This share dropped to 1 percent in March 2016. State-level data from the U.S. Bureau of Labor Statistics show that the unemployment rate in Texas has nevertheless remained low; in fact, it fell from 5.1 percent in June 2014 to 4.3 percent in March 2016, which is below the national average. This means that, although one in five workers in the mining and logging sector lost their job, most of these unemployed workers found employment in other sectors in Texas (or must have relocated to other states, presumably for new jobs there). The fact that the Texan economy apparently was able to absorb most of these 70,000 workers among the pool of close to 12 million employed, while the Texan labor force increased by 2.1 percent (270,850 workers) at the same time

8. In related work, Feyrer, Mansur, and Sacerdote (2015) conclude based on estimates of county-level regressions that the shale boom created 725,000 jobs (two-thirds of which are in the mining sector), which they equate with a reduction in the U.S. unemployment rate of 0.5 percentage point during the Great Recession. These estimates, however, combine job gains from shale oil as well as shale gas production, and they do not allow for the possibility that job gains near shale counties may coincide with job losses elsewhere. Leaving aside these caveats, it is clear that even a partial reversal of these job gains presumably would have resulted in an increase in the unemployment rate of several percentage points, if frictional unemployment were empirically important. What figure 8 shows is that the U.S. unemployment rate continued to fall at a steady rate from 6.1 percent in mid-2014 to 5 percent in March 2016, rather than increasing relative to the previous trend.

(consistent with the view that Texas may have absorbed oil workers returning from other states as well), speaks against the existence of important frictions preventing the reallocation of labor. Of course, this point is difficult to verify, given that there are other reasons for labor migration. What matters for our purposes is that the decline in the unemployment rate is not a statistical artifact of a higher labor force, given that the number of unemployed decreased by 12.2 percent, while the number of employed increased by 2.8 percent. In short, the change in the unemployment rate since June 2014 appears inconsistent with large multiplier effects from the oil sector to other sectors of the Texan economy, at least at the 21-month horizon.[9]

Even in a state such as North Dakota—where, as of June 2014, 6.4 percent of all jobs were in the mining and logging sector, and where almost every second worker in this sector lost his or her job—the unemployment rate rose only slightly, from 2.7 to 3.1 percent. A natural conjecture is that this performance was made possible by the migration of unemployed workers to other states. If this interpretation were correct, one would expect a decline in the civilian labor force split between a decline in the number of the employed and in the number of the unemployed such that the unemployment rate, defined as the number of unemployed residents divided by the labor force, remains approximately stable. As it turns out, the data suggest a different pattern. North Dakota has actually experienced an increase in its labor force and in the number of unemployed since June 2014, accompanied by a decline in the number of employed. The latter decline has been surprisingly modest (–0.1 percent), despite substantial job losses in the nonfarm sector (–4.6 percent), and in mining and logging in particular (–41.1 percent). Moreover, the substantial increase in the number of unemployed in North Dakota (+18.6 percent, starting from a small base) has been partially offset by an increase in the civilian labor force (+0.4 percent, starting from a much larger base), which explains the modest increase in the unemployment rate. One interpretation of this evidence is that natural population growth and, possibly, continued migration into North Dakota after June 2014, explain the increase in the number of unemployed residents and the increase in the civilian labor force, as well as the remarkable stability of the unemployment rate.

9. An interesting question is how these former oil workers have been absorbed by the economy. Herrera, Karaki, and Rangaraju (2016) provide evidence for the reallocation of jobs lost in the oil and natural gas sector to the service sector, manufacturing sector, and construction sector.

Table 5. Changes in Labor Market Indicators in U.S. Oil-Producing States, June 2014–March 2016

State	*Labor force*	*Number of employed*	*Number of unemployed*	*Unemployment rate (percent)*	*Percent share of mining and logging jobs in employment*
Alaska	−4,900	−3,200	−1,700	−0.4	−0.4
Montana	9,500	10,900	−1,500	−0.3	−0.5
New Mexico	−1,000	4,000	−5,100	−0.6	−1.0
North Dakota	1,700	−400	2,100	0.4	−2.5
Oklahoma	82,700	80,700	2,100	−0.1	−0.9
Texas	270,600	351,100	−80,600	−0.8	−0.7
Wyoming	−6,000	−8,800	2,800	1.0	−2.1

Source: U.S. Bureau of Labor Statistics.

Table 5 summarizes the evidence for all seven "oil states" in the United States (defined as states with an oil share in value added above 5 percent, as discussed in more detail in subsection V.A). This evidence suggests three main conclusions. First, between June 2014 and March 2016 all seven oil states experienced declines in the share of jobs in mining and logging. These declines ranged from 0.4 to 2.5 percentage points. Second, nevertheless, the overall unemployment rate declined in all but two of these oil states, and in the latter two states the increase in the unemployment rate was quite small. Third, only in Alaska and Wyoming is there evidence of the unemployment rate being stabilized by the unemployed as well as formerly employed workers leaving the state. In contrast, four of the seven oil states experienced an increase in the labor force, often associated with a strong increase in employment, as in the case of Montana, Texas, and Oklahoma. New Mexico, in contrast, saw little change in its labor force, but a large reduction in the number of its unemployed. We conclude that unemployment, whether voluntary or not, has remained remarkably low in the oil states, providing evidence against a quantitatively important reallocation effect, at least in the oil sector. There is little evidence that unemployed workers waiting out the slump have been driving up the unemployment rate in these oil states, unlike what one might have expected based on Hamilton's (1988) model.[10]

10. Given that the U.S. shale oil industry is capital intensive, one may ask what the evidence is that capital embodied in oil machinery has been underutilized, following the oil price decline. Measuring the underutilization in capital is not straightforward, as mentioned above, but there is some suggestive evidence. For example, the number of oil rigs

III.B. Did Uncertainty about Future Gasoline Prices Hold Spending Back?

The evidence in subsection III.A casts doubt on the notion that severe frictions impeding the reallocation of labor and capital have been responsible for an economic slowdown that offset the stimulus computed in section II; but there is an alternative potential explanation for the U.S. economy's weak response to lower gasoline prices that focuses on a different channel. This alternative explanation postulates that an increase in uncertainty about future oil and gasoline prices may be responsible for holding back consumption and investment spending, and hence real GDP growth.

In this subsection, we focus on the question of whether increased uncertainty about the future price of gasoline may have partially offset the discretionary income and operating cost effect documented in section II. In particular, increased gasoline price uncertainty could be the reason why consumers chose not to buy more automobiles, helping to explain why the consumption of new motor vehicles fell relative to trend, as shown in table 1, at a time when gasoline prices were lower than they had been for a long time. The argument is that the decision to buy a new vehicle depends in part on consumers' expectations of future gasoline prices. If future gasoline prices become more uncertain, it makes sense for consumers to hold off buying a new car for the time being, even when expected gasoline prices are low. This point is closely related to Bernanke's (1983) model of how increased uncertainty about the price of oil may cause a delay in investment projects.[11] The same reasoning applies to purchases of consumer durables such as cars and light trucks. The quantitative importance

---

in the United States, as reported by Baker Hughes, has declined by about 75 percent since its peak in October 2014, suggesting considerable underemployment of capital embodied in rigs. Likewise, rail traffic data reported by the Association of American Railroads show that the average weekly number of carloads of petroleum and petroleum products has declined by more than 30 percent since its peak in September 2014, suggesting ample underutilization of the fleet of tanker rail cars. This problem is not limited to the oil sector. One would expect the underutilization of capital to extend more broadly to other sectors of the economy in all those states where oil is produced. In subsection V.A, we return to this question. We quantify the extent to which reduced economic growth in these oil-producing states has affected U.S. real GDP growth and show that the underutilization of capital in the oil-producing states had at most a very small effect on the overall growth of the U.S. economy.

11. Bernanke's (1983) point is that—to the extent that the cash flow from an irreversible investment project depends on the price of oil—real options theory implies that, all else equal, increased uncertainty about the real price of oil prompts firms to delay investments, causing investment expenditures to drop. Related work includes Pindyck (1991).

of this effect depends on how important the real price of gasoline is for automobile purchase decisions and on the share of such expenditures in aggregate spending.[12]

To assess the empirical content of the real options model, we must construct a measure of consumers' uncertainty about future gasoline prices at the horizons relevant to purchases of automobiles. One challenge is how to measure uncertainty at the horizons longer than the usual monthly or quarterly horizon. The other challenge is that we are concerned with the uncertainty perceived by consumers rather than by financial markets (as embodied in options prices). Similarly, commonly used measures of price uncertainty based on the conditional variance in generalized autoregressive conditionally heteroskedastic (GARCH) models need not be good proxies for the uncertainty of U.S. consumers. In addition, GARCH estimates are backward-looking by construction, and extrapolating from monthly or quarterly GARCH models to multiyear horizons is inherently problematic. We therefore consider an alternative proxy for gasoline price uncertainty, defined as the standard deviation of the responses of participants in the Michigan Survey of Consumers to the question about the expected change in the price of gasoline at the one-year and the five-year horizons.[13]

Figure 9 suggests a pronounced increase in consumers' uncertainty about gasoline prices, at both short and longer horizons, in late 2014. Note that not all increases in gasoline price uncertainty are exogenous with respect to U.S. consumption. For example, the tremendous surge in gasoline price uncertainty in 2008 and 2009 was clearly driven by the recession associated with the financial crisis. The spike in gasoline price uncertainty after June 2014, in contrast, was not caused by a U.S. recession; hence, for our purposes, it may be viewed as a potential explanation for consumers' purchases of motor vehicles.

12. A closely related argument is made by Edelstein and Kilian (2009), who observe that increased uncertainty about the prospects of staying employed in the wake of unexpected changes in the real price of oil may cause an increase in precautionary savings (or, equivalently, a reduction in consumer expenditures). In this interpretation, uncertainty about gasoline prices may affect not merely consumer durables such as cars that are fuel-intensive in use, but other consumer expenditures as well. Here we focus on the uncertainty effect on the consumption of motor vehicles.

13. Disagreement among individual survey respondents' predictions is not in general the same as any one respondent's uncertainty about future outcomes, but Zarnowitz and Lambros (1987) provide evidence, in the context of inflation expectations, that the standard deviation of the responses across respondents and the standard deviation of individual predictive distributions tend to be positively correlated, especially at lower frequency. For related evidence in a different context, see Bachmann, Elstner, and Sims (2013).

Figure 9. Measures of U.S. Consumers' Uncertainty about the Future Price of Gasoline, January 2006–March 2016[a]

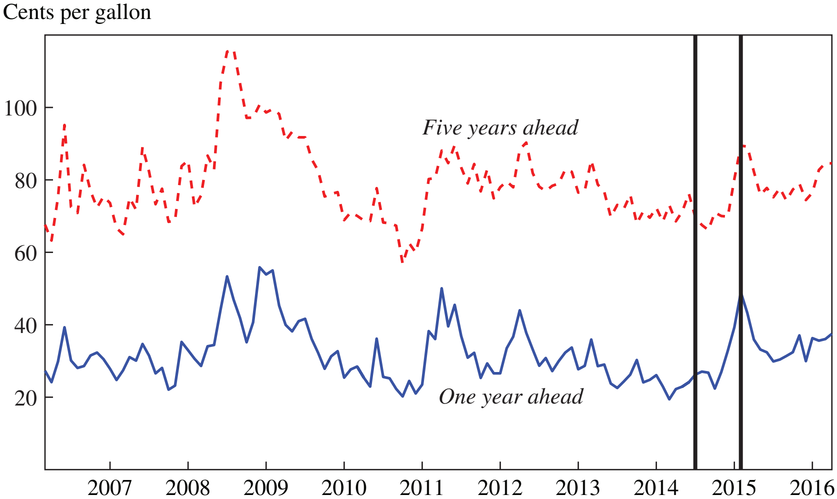

Source: Michigan Survey of Consumers (courtesy of Richard Curtin).

a. Uncertainty is measured by the standard deviation of the responses of survey participants to the question about the expected change in the price of gasoline one year and five years ahead. The vertical lines are at June 2014 and January 2015, when uncertainty peaked.

The literature on the uncertainty effect suggests that this spike in uncertainty, all else equal, should have been associated with a reduction in vehicle sales. Indeed, the upper panel of figure 10 shows that U.S. sales of autos and light trucks remained sluggish between June 2014 and January 2015, before accelerating in the second half of 2015. This evidence would seem to be supportive of a quantitatively important uncertainty effect, except for the fact that current conditions for buying a vehicle, as recorded by the Michigan Survey of Consumers, greatly improved in late 2014, directly contradicting this hypothesis. If consumers chose not to buy a new car despite the strong improvement in current buying conditions, then the reason cannot have been higher gasoline price uncertainty, but must have been some other economic development that offset the stimulative effect of lower gasoline prices, adding credence to the standard linear model of the transmission of purchasing power shocks.

This conclusion is reinforced by the fact that there is clear evidence of substitution across classes of vehicles with different levels of fuel

Figure 10. Sales and Current Buying Conditions for Vehicles, May 2014–March 2016[a]

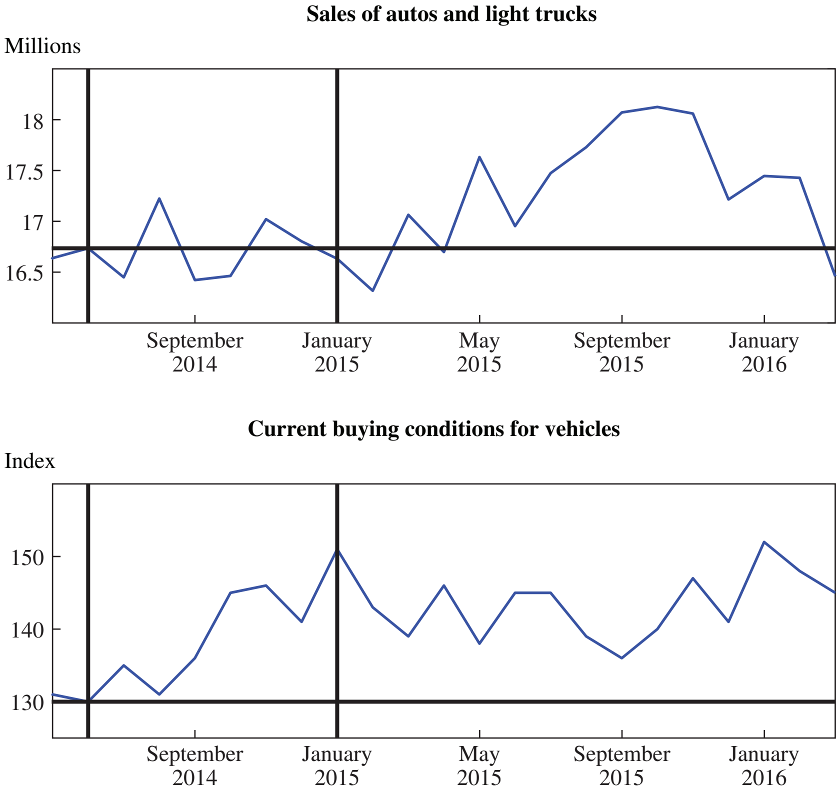

Sources: U.S. Bureau of Economic Analysis; Michigan Survey of Consumers.
a. The vertical lines are at June 2014 and January 2015, when uncertainty peaked.

efficiency. If consumers choose to buy a light truck rather than a car, for example, this fact indicates that they are not deterred by gasoline price uncertainty, but are quite confident in buying a type of vehicle that is clearly less fuel efficient than the alternatives. The top panel of figure 11 shows that after June 2014, auto sales actually declined, while sales of light trucks increased faster than overall vehicle sales, providing additional evidence against an important role for gasoline price uncertainty. The share of light trucks in total light vehicle sales increased from 53 percent in June 2014 to 59 percent in March 2016. The bottom panel of figure 11 shows that there has been a disproportionate decline in the sales

Figure 11. Decomposition of Vehicle Sales, January 2010–March 2016[a]

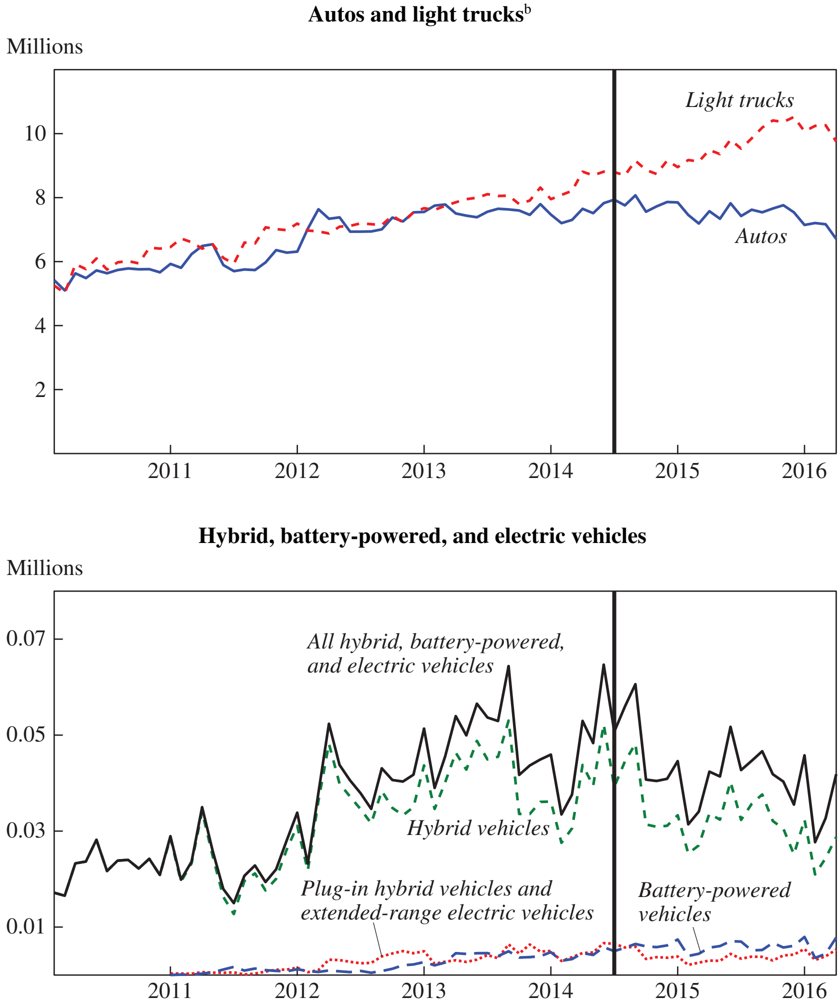

Sources: U.S. Bureau of Economic Analysis; Electric Drive Transportation Association.
a. The vertical lines mark June 2014, the month before the oil price decline began.
b. Aggregate of domestic and foreign sales.

of hybrid cars since June 2014 relative to overall auto sales, corroborating our earlier evidence.[14]

IV. How Much Did the Unexpected Oil Price Decline Stimulate Nonresidential Investment, Excluding the Oil Sector?

Another form of private spending that may be stimulated by unexpectedly low oil and fuel prices is private nonresidential investment. In this section, we focus on private nonresidential investment, excluding the oil sector. The response of oil-related investment to unexpectedly low oil prices is analyzed in section V. There are two primary channels through which unexpectedly low oil prices may stimulate nonresidential investment not related to oil. One channel is that firms directly benefit from lower fuel prices to the extent that they purchase fuel and equipment that uses fuel. This channel is not likely to be quantitatively important outside the transportation sector. The other channel is that, with lower oil prices lifting household income, higher consumer spending encourages business capital spending more broadly (Yellen 2011).[15]

Let $\Delta inv_t^{ex\,oil}$ denote the quarterly growth rate of real private nonresidential investment (excluding structures and equipment investment by the oil sector), demeaned to account for the change in average investment growth after December 2008. Given that the magnitude of the nonresidential investment stimulus largely depends on the consumption stimulus, we allow $\Delta inv_t^{ex\,oil}$ to depend linearly on the same purchasing power shock measure as in the baseline consumption model, suitably aggregated to quarterly frequency:

$$\Delta inv_t^{ex\,oil} = \sum_{i=1}^{4} \beta_i \Delta inv_{t-i}^{ex\,oil} + \sum_{i=0}^{4} \gamma_i PP_{t-i} + u_t, \tag{2}$$

14. In contrast to the sales of hybrid vehicles, the sales of battery-powered vehicles have not responded to the decline in the price of gasoline, suggesting that buyers of electric cars are primarily motivated by environmental concerns and less by fuel costs.

15. These effects may be offset by higher oil and gasoline price uncertainty, to the extent that the cash flow from investments depends on the prices of oil and gasoline. For example, Kellogg (2014) documents that higher oil price uncertainty affected the investment decisions made by oil producers in Texas. Given that our analysis in this section excludes the oil sector, this uncertainty effect may be safely ignored. Not only is the price of fuel not an important determinant of the cash flow of most nonresidential investment projects in the economy, but, in addition, we have already established that even for automobile purchases, where this effect should be most pronounced, the uncertainty channel of transmission does not seem empirically relevant.

where u_t denotes the regression error.[16] The estimated cumulative stimulus for $inv_t^{ex\ oil}$ between 2014:Q2 and 2016:Q3 is 2.2 percent. Given the share of 11.8 percent of non-oil private nonresidential investment in U.S. real GDP in 2014:Q2, this implies a cumulative increase in real GDP of 0.22 percent after accounting for an import propensity of 0.15. We also estimated an alternative model that allows PP_t to incorporate changes in the U.S. economy's dependence on oil and gasoline imports, as discussed in section II. The implied cumulative stimulus from this alternative model is 0.19 percent, which is almost identical to the baseline estimate.

V. Why This Time Might Be Different

Even under the maintained assumption of a linear relationship between purchasing power shocks and real consumption growth (or real non-oil nonresidential investment growth), we need to consider the possibility that the transmission of this latest oil price shock may be different because of latent structural changes in the U.S. economy. One potential source of such temporal instability is the increased importance of the shale oil sector for the U.S. economy after 2011.

V.A. How Important Has the Contribution of the Shale Oil Sector Been to U.S. Real GDP?

By mid-2014, U.S. shale oil production alone accounted for one quarter of all crude oil used by the U.S. economy (Kilian 2016b). A view that has gained popularity is that low oil prices may be harmful to the U.S. economy because of their disruptive effects on the domestic oil industry, notably the shale oil sector.

THE EFFECTS OF THE SHALE OIL SECTOR ON VALUE ADDED One way of assessing the empirical content of this proposition is to quantify the reduction in the value added generated by the oil industry following the decline in the price of oil since June 2014. Although U.S. oil production initially continued to increase, reflecting substantial productivity increases in extracting shale oil, and peaked only in April 2015, the U.S. oil sector experienced a severe contraction in 2015–16. This contraction is clearly visible in measures of gross output such as the number of barrels of crude oil produced by the industry as well as in data on employment and capital expenditures.

16. The lag order matches that used in Edelstein and Kilian (2007). The sample period is 1970:Q3 to 2016:Q1.

Table 6. Actual and Counterfactual Average U.S. Real GDP Growth[a]

Measure of real GDP growth	*2014:Q3–2016:Q1*	*2014:Q3–2015:Q2*	*2015:Q3–2016:Q1*
Real GDP (value added)[b]	2.38	2.72	1.70
Excluding mining sector	2.43	2.55	2.21
Mining sector	2.45	9.06	–9.59
Real GDP[b,c]	2.38	2.69	1.76
Excluding oil-producing states	2.33	2.55	1.91
Oil-producing states	2.70	3.71	0.72
Real GDP	2.19	2.72	1.48
Excluding the change in the petroleum trade balance	2.16	2.69	1.46
Excluding the change in investment in the oil sector	2.56	3.06	1.90

Sources: U.S. Bureau of Economic Analysis; authors' calculations.

a. Growth is measured in percentage changes at annual rates.

b. Sample ends in 2015:Q4.

c. The state-level counterfactual is based on real GDP as reported in the regional economic accounts, and differs slightly from real GDP as reported in the national income and product accounts.

Assessing the magnitude of the effect of this contraction on real value added is not straightforward because there are no quarterly value added data on U.S. shale oil production (or, for that matter, on total oil production). The closest available aggregate is mining, which includes oil and natural gas extraction, other mining activities, and support services for all mining activities. The top panel of table 6 shows that the overall effect of changes in mining on real GDP growth between 2014:Q2 and 2015:Q4 has been negligible. This result obscures the fact that between 2014:Q2 and 2015:Q2, growth in mining value added actually raised U.S. real GDP growth by 0.17 percentage point at annual rates, whereas after 2015:Q2 it lowered real GDP growth by 0.51 percentage point at annual rates, as value added in mining fell by 9.6 percent. Further inspection of the annual real value added data, which provide a more detailed breakdown, suggests that oil and natural gas extraction combined, far from contracting, actually continued to grow even in 2015 at an astounding rate of 16 percent, even as other mining activities and overall mining support declined by 7 and 14 percent, respectively. This evidence suggests that much of the contraction in the shale oil industry occurred not in production but in support services. The reason why these changes do not matter more at the aggregate level is not only that some of the changes are offsetting, but also that the share of mining in GDP has remained quite small, having risen gradually from 2.2 percent in 2007 (before the shale oil boom) to a peak of 2.6 percent in 2013, before falling to 1.7 percent in 2015.

Focusing on the direct contribution of the oil sector may be underestimating its overall impact on value added, however. Clearly, oil states such as North Dakota and Texas experienced an economic boom between 2010 and 2015 that extended to the service sector, residential housing, and other infrastructure required to sustain higher levels of oil production (Feyrer, Mansur, and Sacerdote 2015). When the price of oil fell and the boom turned into a bust, many other sectors of the economy in the oil-producing states also contracted. It is difficult to measure these effects directly, but a simple thought experiment allows us to bound these broader impacts at the state level on U.S. real GDP. The BEA provides data on real GDP growth for every U.S. state. We classify these states into those with an oil share in their value added in 2014 above 5 percent (referred to, as noted above, as the oil states) and those with a lower share. The oil states include North Dakota (with an oil share in value added of 84 percent), Alaska (40 percent), Wyoming (21 percent), New Mexico (14 percent), Texas (8 percent), Oklahoma (7 percent), and Montana (6 percent).[17] These states also include the most important shale oil sites in the country (Kilian 2016a). We then ask how different U.S. real GDP growth would have been if these oil states had not been part of the U.S. economy. This approach allows us to capture both the direct effects and the indirect state-level effects of the decline in shale oil production on U.S. real GDP growth.

The middle panel of table 6 shows that after excluding the seven oil states from the U.S. economy, the aggregate rate of growth would have been only marginally different, suggesting that the state-level effects of the decline in shale oil production on value added are quite modest. In fact, between 2014:Q2 and 2015:Q4, shale oil states overall slightly increased U.S. real GDP growth from 2.3 to 2.4 percent at annual rates rather than lowering it. Only starting in 2015:Q3, when growth in the oil states had dropped from 3.7 to 0.7 percent at annual rates, is there any evidence that these states pulled down aggregate real GDP growth. The counterfactual growth rate exceeded the actual growth rate by 0.15 percentage point. This evidence suggests that if the shale oil sector was indeed responsible for the

17. Following Hamilton and Owyang (2012), the state-level oil share is calculated as 100 times the number of barrels of crude oil produced in a given state in 2014, as reported by the EIA, weighted with the annual domestic first purchase price (dollars per barrel) for that year, and then divided by the 2014 state personal income, as reported by the BEA.

Table 7. Actual and Counterfactual Average U.S. Real Investment Growth[a]

Measure of real investment growth	*2014:Q3–2016:Q1*	*2014:Q3–2015:Q2*	*2015:Q3–2016:Q1*
Private fixed nonresidential investment[b]	1.5	3.8	−1.4
Excluding oil-related investment	4.6	6.8	1.8
Oil-related investment only	−48.2	−35.2	−61.5
Investment in structures	−2.9	0.2	−6.7
Excluding petroleum and natural gas structures	10.2	12.0	7.9
Petroleum and natural gas structures only	−50.8	−34.2	−66.6
Investment in equipment	1.6	3.3	−0.6
Excluding mining and oil field machinery	2.7	4.7	0.2
Mining and oil field machinery only	−39.4	−39.7	−39.1

Sources: U.S. Bureau of Economic Analysis; authors' calculations.

a. Growth is measured in percentage changes at annual rates.

b. Oil-related investment includes petroleum and natural gas structures as well as mining and oil field machinery.

sluggish growth of the U.S. economy, there must have been other transmission channels at play. There are several such mechanisms to consider.

THE EFFECTS OF SHALE OIL ON REAL GDP THROUGH FIRMS' INVESTMENT EXPENDITURES To the extent that variation in the growth rate of real GDP is disproportionately affected by variation in the growth rate of real investment, it is conceivable that the oil sector may have had large effects on economic growth without having a large direct effect on value added. It is widely accepted that the unprecedented expansion of the U.S. shale oil sector has been a major contributor to aggregate investment since 2010, changing the dynamics of the U.S. economy. As a result, when the price of oil fell after June 2014, real investment in the U.S. oil sector dropped sharply, which could help explain why U.S. aggregate real nonresidential investment did not expand nearly as much in response to lower oil prices, as one might have expected.

The bottom panel of table 6 and table 7 examine the quantitative importance of this effect. Investment in the oil sector is approximated by the sum of investment in mining and oil field machinery and investment in petroleum and natural gas structures. The top panel of table 7 shows that total real fixed nonresidential investment in the U.S. economy between 2014:Q2 and 2016:Q1 on average increased by 1.5 percent at annual rates, compared with 2.2 percent growth in real GDP. Over the same period, oil investment dropped at an annual rate of 48.2 percent. Thus, after excluding investment in the U.S. oil sector, real private fixed nonresidential investment would have increased at a rate of 4.6 percent, about three times as fast as the actual data. The middle panel of table 7 shows that investment in structures would have grown at a rate of 10.2 percent (rather than declining at a rate of 2.9 percent), and the bottom panel shows that investment in equipment would have grown at a rate of 2.7 percent (rather than merely 1.6 percent). This robust growth was largely offset by reduced investment in the oil sector, however. This mechanism is not new. It has already been documented by Edelstein and Kilian (2007) in the context of the 1986 oil price decline. What is new is the magnitude of the decline in the real price of oil, on one hand, which was twice as large after 2014:Q2 compared with 1985:Q4, and the magnitude of the decline in oil-related investment, on the other hand, which amounted to 48 percent between 2014:Q2 and 2016:Q1 compared with only 21 percent between 1985:Q4 and 1987:Q3. Given that the share of oil and natural gas extraction in GDP was 1.7 percent in 1985 as well as in 2014, a likely explanation for the disproportionate drop in oil investment is that the price of oil in 2014–16 declined by about twice as much.

A complementary explanation could be that investment by shale oil producers is more price sensitive than investment by conventional oil producers. Whether this common perception is actually correct is not clear. The decision to continue to invest in shale oil production depends on whether the expected price of oil exceeds the long-run marginal cost of oil production. If so, oil production remains profitable and investment continues. Otherwise, investment ceases. One difference from conventional oil production is that the marginal cost of producing shale oil tends to be higher than that for conventional oil production—which, all else equal, suggests that, as the expected price of oil declines, investment by shale oil producers should cease before conventional oil investment. Another difference, however, is that investment in the shale oil sector has a much shorter horizon. Thus, the decision to cut shale oil investment only depends on the expected evolution of the price of oil in the short run. For conventional investment, in contrast, the price of oil expected at longer horizons also matters. For example, expectations of a longer-term price recovery would tend to make conventional oil investment more robust to oil price declines than shale oil investment. Which type of investment is affected more is therefore ambiguous, in general. In addition, it needs to be kept in mind that the uncertainty about the future price of oil may be higher in the short run than in the longer run, which would slow investment in shale oil compared with longer-term oil investments. If oil price uncertainty is lower in the short run than in the long run, in contrast, shale oil investment would be boosted relative to investment in conventional oil. Thus, it is not clear a priori whether shale oil investment is more responsive to oil price fluctuations than other types of oil investment.[18]

The bottom panel of table 6 shows the effect of reduced oil investment on U.S. real GDP growth. The decline in average real GDP growth associated with lower oil prices is less dramatic than the estimates in table 7, reflecting the comparatively low share of total investment in GDP compared with the share of consumption in GDP, but is still economically significant. U.S. real GDP would have increased at an average annual rate of 2.6 percent, excluding the decline in investment in the oil sector, compared with the 2.2 percent observed in the data. Thus, lower oil-related investment accounts for a reduction of 0.4 percentage point in U.S. real GDP growth measured at annual rates.

THE EFFECTS OF SHALE OIL ON REAL GDP THROUGH THE PETROLEUM TRADE BALANCE Lower oil prices may affect real GDP by changing consumption and investment expenditures, but also by changing net petroleum exports. As long as the volume of oil imports remains unchanged, a change in the real price of oil leaves real oil imports unchanged. If a lower real price of oil discourages domestic oil production, however, for given U.S. oil consumption, real oil imports must increase. This effect (which mirrors the changes in value added by the oil sector) must be included in modeling the effects of lower real oil prices on the expenditure side of real GDP. The analysis is further complicated by the rise of the U.S. shale oil sector, however.

In quantifying the effect of lower oil prices on the trade balance after 2014:Q2, it makes sense to focus on the petroleum trade balance of the U.S. economy rather than the crude oil trade balance, where petroleum is defined to include both crude oil and refined products. The reason is that the U.S. shale oil revolution not only permitted U.S. refiners to curtail their oil imports but also allowed refiners to export refined products such as gasoline and diesel fuel on a much larger scale than heretofore (Kilian 2016a, 2016b). Although U.S. net petroleum imports during the last seven years have fallen from $240 billion to $102 billion (in chained 2009 dollars), the United States has remained a net petroleum importer. However,

18. It has been argued that the reduction in the oil sector's real investment may have caused real investment in other sectors of the economy to decline as well. If so, one would expect a similarly sharp drop in these components of investment after 2014:Q2. Time series plots of investment in industrial equipment and investment in transportation equipment, however, are not supportive of such a link. Only railroad equipment investment mirrors the decline in oil investment, but it can be shown that this pattern primarily reflects a decline in traffic volumes in other commodities such as coal rather than reductions in petroleum shipments.

contrary to the conventional argument above, the U.S. petroleum trade balance actually improved, following the 2014–16 decline in the price of oil, with exports growing faster than imports, driven by continued shale oil production. The bottom panel of table 6 shows that excluding the change in the petroleum trade balance since 2014:Q2 from real GDP would have slightly lowered average real GDP growth by 0.03 percentage point at annual rates. This improvement in the petroleum trade balance of the U.S. economy contributed to real GDP growth, reinforcing the consumption and investment stimulus discussed in sections II and IV.

THE EFFECTS OF THE SHALE OIL SECTOR ON REAL GDP THROUGH FINANCIAL SPILLOVERS Another channel through which the decline in the price of oil may slow down the economy is by exposing banks and other financial institutions to oil price risks. Following the financial crisis, bank lending to shale oil producers was considered a growth market that offered high returns at seemingly low risk. Banks actively sought to finance both large and small oil companies without much regard for these companies' cash flows. In many cases, oil below the ground was considered sufficient collateral. Because the price of oil underpins the value of the assets securing these loans, the decline in the price of oil after June 2014 increased banks' oil exposure. At the same time, lower oil prices reduced the cash flow generated by oil producers, making it more difficult for borrowers to service their loans and raising the probability of defaults. Moreover, as the price of oil fell, debt-ridden producers had an incentive to increase output to cover interest payments, in turn putting further downward pressure on the price of oil.

By late 2015, there was growing concern about bank reserves proving inadequate to deal with nonperforming loans to the oil sector, about preapproved unsecured credit lines to oil and gas companies, and about banks being subject to additional undisclosed oil price risks. By the early and middle parts of 2016, many major banks in turn were attempting to quell concerns about bad oil loans by raising reserves and by disclosing likely losses. These concerns arose despite the fact that bank loans to oil and gas companies account for at most 5 percent of total loans at the major U.S. banks and in many cases for far less, making these banks' exposure much lower than their exposure to mortgage risk before the U.S. housing crisis.

Figure 12 plots a stock market index designed to track the performance of 24 U.S. bank stocks. It shows that bank stocks initially appreciated amid falling oil prices. As the number of bankruptcies in the oil and gas extraction sector increased and the banks' oil exposure became more widely known, the values of these banks' stocks fell sharply, reaching a trough in January 2016. Their partial recovery starting in February closely tracks the

Figure 12. KBW Bank Index, January 2009–May 2016[a]

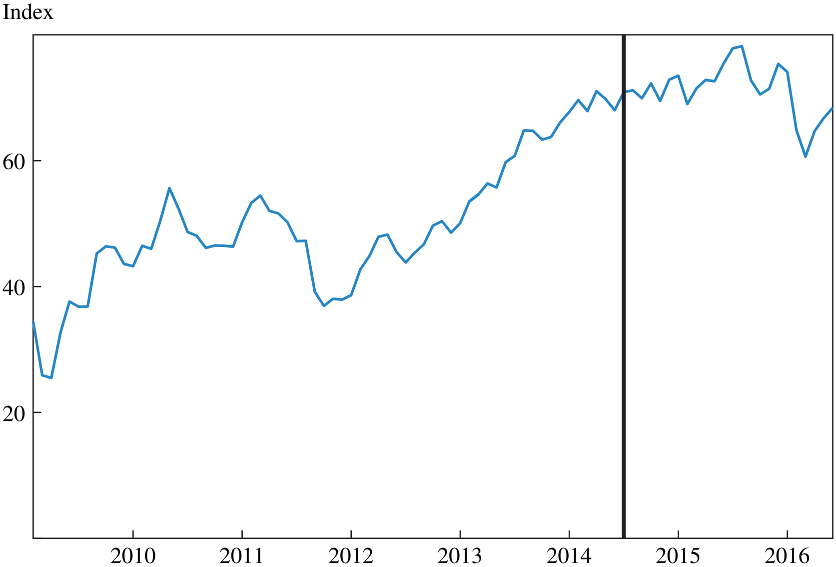

Source: Keefe, Bruyette & Woods.

a. The KBW index includes 24 U.S. banking stocks, and is traded on the Philadelphia Stock Exchange. The index has been expressed as a monthly average. The vertical line is at June 2014, the month before the oil price decline began.

partial recovery of the price of crude oil, which helped alleviate concerns about the ability of oil producers to service their debt and about the diminishing value of the banks' collateral. Yet overall, there is no evidence that financial fragility has been a cause of the economic slowdown that started in early 2015. In fact, at that point in time, bank stocks were still appreciating. Nor is there evidence that the growing number of bankruptcies in the oil and natural gas extraction sector has been spreading to other sectors.[19]

19. BankruptctyData.com collects monthly information on corporate bankruptcies based on daily court filings. These data show that there has been a strong increase in the number of bankruptcies in the oil and natural gas extraction sector (SIC code 13), from 0 bankruptcies in June 2014 to 82 in May 2016, reaching a cumulative total of 560. This sector includes crude oil, petroleum, and natural gas producers; firms involved in drilling oil and gas wells; oil and gas exploration services; and other oil and gas field services. A detailed analysis (not reported here to conserve space) shows that, among the 74 remaining two-digit SIC industries, there is not one industry that exhibits an increase in the number of bankruptcies that resembles that of the oil and gas extraction industry. Thus, we conclude that this channel has not been quantitatively important so far, although it may yet contribute to an economic slowdown.

V.B. Has There Been a Shift in Consumers' Savings Behavior?

As our back-of-the-envelope calculation showed, consumers' purchasing power increased after June 2014. If consumers did not spend this extra income, as presumed by conventional economic models, where did this income go? One possibility is that consumers took the opportunity to pay off mortgage or credit card debt, to increase their savings, or to acquire financial assets on a scale not seen in historical data. Such an unprecedented shift in consumers' savings behavior after June 2014 would invalidate the predictions of the linear model of the transmission of purchasing power shocks for real private consumption.

There is no empirical support for this view, however. BEA data show that the personal savings rate of U.S. households, defined as after-tax disposable income minus personal outlays as a percentage of after-tax disposable income, actually slightly declined, from 5.9 percent on average between January 2009 and June 2014 (after excluding an outlier in November and December 2012 associated with changes in fiscal policy) to 5.8 percent on average between July 2014 and March 2016. In fact, from June 2014 to March 2015, when the bulk of the oil price decline occurred, the savings rate dropped from 5.8 percent to as low as 5.3 percent at one point, before recovering later in 2015. Only between August 2015 and March 2016 did the savings rate exceed its long-run average, reaching 6.2 percent in March 2016. The increment in the savings rate of 0.3 percentage point relative to June 2014 is much smaller than the increment of about 1 percentage point that one would have expected, all else equal, if the cumulative gain in discretionary income since June 2014 had been entirely converted into savings. Likewise, flow-of-funds data from the Federal Reserve (which are not shown, to conserve space), provide no support for the deleveraging hypothesis. Households increased their liabilities, in some cases at an increasing rate, rather than reducing them.

VI. What Is the Net Stimulus?

The increased importance of shale oil documented in section V complicates the assessment of the U.S. economy's overall response to lower oil prices. Table 8 summarizes the cumulative effects on aggregate spending that we have identified thus far from a national income accounting point of view, focusing on the three components of the national income identity $GDP = C + I + G + X - M$ that are directly affected by the oil price decline, where C denotes private consumption, I denotes private investment, G denotes government spending, and $X - M$ denotes the external balance.

Table 8. The Net Stimulus from Unexpectedly Low Real U.S. Oil Prices[a]

	Percent of cumulative real GDP growth	
Effect on real GDP of	*2014:Q3–2016:Q1*	*1986:Q1–1987:Q3*
Discretionary income effect on private consumption	0.61	0.28
Operating cost effect on private consumption	0.09	0.08
Oil-related private nonresidential investment	–0.57	–0.43
Non-oil-related private nonresidential investment	0.22	0.11
Petroleum trade balance	0.04	–0.41
Net stimulus	0.39	–0.37

Source: Authors' calculations.

a. The estimates have been adjusted based on a marginal import propensity of 0.15, and take into account the share of each expenditure component in real GDP. A net stimulus of 0.39 percent translates to an increase in the average growth rate of real GDP of 0.2 percentage point at annual rates.

The baseline model of private real consumption presented in section II implied that the discretionary income effect cumulatively raised real consumption by 1.05 percent. Weighting this result with the share of consumption in GDP of about 69 percent and adjusting it for a marginal propensity to import of 0.15, we obtain a stimulus to cumulative real GDP growth of 0.61 percent. The corresponding operating cost effect adds another 0.09 percentage point after accounting for higher imports. The stimulus arising from non-oil-related nonresidential investment is 2.2 percent, which, when weighted by the share of 11.8 percent in GDP and adjusted for an import propensity of 0.15, yields a cumulative increase in real GDP of 0.22 percent. The combined stimulus of 0.9 percent must be traded off against the reduction in cumulative real GDP growth caused by lower real investment in the oil sector broadly defined. In section V, we showed that cumulative real GDP growth fell by 0.67 percent, as oil investment contracted, which reduces to 0.57 after accounting for the implied reduction in imports. Finally, we need to account for the improvement in the petroleum trade balance, as discussed in section V, which raises real GDP growth by 0.04 percentage point cumulatively. This simple exercise implies a net stimulus of 0.39 percent in cumulative real GDP growth (or an increase of 0.2 percentage point in average real GDP growth, at annual rates), which is close to zero.[20] Thus, the fact that

20. Using the alternative model that explicitly allows for changes in the dependence of the U.S. economy on oil and gasoline imports, the net stimulus is 0.2 percentage point, corresponding to an increase of 0.1 percent in the average growth of real GDP at annual rates, which is quite close to the baseline model specification.

the average U.S. real GDP growth, as shown in table 1, accelerated only slightly, from 1.8 percent at annual rates to 2.2 percent, is not surprising.

Another reason why real GDP growth did not increase faster after 2014:Q2 undoubtedly was the slow growth of real nonpetroleum exports after June 2014, which dropped from 2.9 percent during 2012:Q1–2014:Q2 to –0.2 percent during 2014:Q3–2016:Q1. Our analysis thus far abstracted from the fact that the decline in the price of oil after June 2014 did not occur all else equal, but was associated at least in part with a global economic slowdown (Baumeister and Kilian 2016b; Kilian 2016b), which in turn slowed U.S. export growth and hence U.S. real GDP growth. It is difficult to quantify this effect without a fully specified model, but the case can be made that real GDP growth after 2014:Q2 would have increased by about 0.3 percentage point to 2.5 percent on average, if U.S. real nonpetroleum exports had continued to grow at an average annual rate of 2.9 percent.

VII. What Has Changed Compared with the Oil Price Decline in 1986?

Given the results in table 8 for 2014:Q4–2016:Q1, what is the evidence that this time is different from what happened following the sustained oil price decline of 1986? The last column of table 8 quantifies the stimulus caused by unexpectedly low oil prices during the 1986:Q1–1987:Q3 period. Overall, there are more similarities than differences. We already showed that the primary reason why real GDP growth remained sluggish after 2014:Q2 was the sharp decline in oil-related investment expenditures. This pattern is not new. A similar decline occurred after the 1986 oil price decline, as first documented by Edelstein and Kilian (2007). As table 8 shows, in the seven quarters after 1985:Q4, lower oil-related investment created a negative stimulus of 0.43 percent of cumulative U.S. real GDP growth, after accounting for the implied change in imports. The more negative stimulus from oil-related investment after June 2014 (0.57 percent) is not unexpected, given that the share of oil and natural gas extraction in GDP was about the same in 2014 as in 1986 and the decline in the price of oil was about twice as large.[21] Table 8 suggests that

21. Although the decline in the price of oil after 2014:Q2 was similar for the first seven months to the price decline after 1985:Q4 (–55 percent versus –57 percent), in 1986–87 the price of oil recovered in the following 14 months (offsetting half the initial decline), whereas in 2015–16 the price of oil (and oil investment) continued to fall even further, with the cumulative decline in the price of oil reaching –66 percent.

oil-related investment since 2014:Q2, if anything, has been more resilient to the decline in the price of oil than in 1986.

The positive stimulus from higher consumer spending (+0.36 percent) and non-oil-related investment spending (+0.11 percent) was about half as large following the 1986 oil price decline, consistent with the cumulative decline in gasoline prices being only half the magnitude of the decline starting in June 2014.[22] The comparatively large operating cost effect of +0.08 percent is driven by the much higher share of new motor vehicles in private consumption during the 1986 oil price decline. The key difference between the two episodes is the negative response of the petroleum trade balance (–0.41 percent) during the 1986:Q1–1987:Q3 period, compared with a slightly positive response (+0.04 percent) in the more recent episode.[23] Even including the latter effect, however, the net stimulus caused by the 1986 oil price decline (–0.37 percentage point) was close to zero, much like the net stimulus in the more recent episode. Thus, the effect of unexpectedly low oil prices on the U.S. economy during the 2014:Q3–2016:Q1 period does not appear fundamentally different from that during the 1986:Q1–1987:Q3 period.

Nor are there large differences in economic performance. Taking account of the decline in the average real GDP growth rate from about 3 percent to 1.9 percent after the financial crisis, in the seven quarters after 1985:Q4, real GDP growth at annual rates was 0.3 percentage point above average; and in the seven quarters after 2014:Q2, it was 0.3 percentage point above average (or 0.6 percentage point, controlling for export growth). Thus, the U.S. economy's overall performance was quite similar in these two episodes, despite the steeper decline in oil-related investment after June 2014. One explanation of this result is that growth in nonresidential investment, excluding oil, dropped in 1986 because of the Tax Reform Act of 1986, whereas there was no such shock in the current episode (Edelstein and

22. These estimates are based on the $PP_t^{alternative}$ measure, given that the dependence of the U.S. economy on oil imports was far below its long-run average in the mid-1980s. The baseline specification implies a somewhat higher stimulus from consumer and non-oil investment spending during the 1986:Q1–1987:Q3 period without affecting the qualitative conclusion that the net stimulus is close to zero.

23. Whereas during the 1985:Q4–1987:Q3 period petroleum exports remained stable and petroleum imports surged to offset lower domestic production, during the 2014:Q2–2016:Q1 period petroleum exports grew faster than petroleum imports, improving the petroleum trade balance and raising real GDP slightly. This outcome was made possible by increased U.S. shale oil production, which facilitated both import substitution and higher petroleum product exports (Kilian 2016a).

Kilian 2007). Controlling for this exogenous event not related to the oil market, real GDP growth after 1985:Q4 should have been higher.

VIII. Conclusions

To summarize, we have shown that the U.S. economy's response to the decline in the real price of oil can be understood based on standard economic models of the transmission of oil price shocks. In particular, we found no evidence that the emergence of the shale oil sector has fundamentally altered the propagation of oil price shocks to the U.S. economy. This fact does not mean that the U.S. shale oil boom did not matter. It is readily apparent that without the shale oil boom, the U.S. economy's response to the recent oil price decline would have been different, if only because of the lower share of oil and natural gas extraction in GDP.

One question of obvious policy interest is whether higher investment in the oil sector could help offset the contractionary effect on private consumption of a future recovery of the real price of oil. The central issue is how fast oil investment would grow in response to an increase in the real price of oil. The argument can be made that new investment in shale oil does not require persistently high expected oil prices. Even a temporary oil price surge would make new investment worthwhile, because shale oil production may respond more quickly to oil price increases than conventional oil production. There are reasons to be cautious about such predictions, however, as emphasized by Kleinberg and others (2016). The rapid expansion of U.S. shale oil production starting in 2009 coincided with the end of the U.S. shale gas boom at the end of 2008. Because shale oil production and shale gas production use the same rigs and hydraulic fracturing equipment, much of the equipment left idle by the shale gas industry was immediately transferred to the shale oil industry, enabling the rapid expansion of U.S. shale oil production. Since October 2014, the rig count has declined by 75 percent. Whether a similar surge in shale oil production could be replicated in response to a higher expected oil price depends on the extent to which the drilling and fracking equipment in question has been scrapped, has rusted away, or has been cannibalized for spare parts since June 2014. The demise of the shale gas sector at the end of 2008 also provided the skilled labor required to operate the equipment. With this labor scattered, following massive layoffs in the shale oil industry, the transition to higher shale oil production in the future is likely to be less smooth than in 2009. In addition, the easy availability of credit played an important role in creating the shale oil boom. It remains to be seen whether shale oil producers will be able to finance new investments as easily as they did in the past.

How an unexpected recovery of the real price of oil would affect U.S. real GDP growth more generally also depends on the determinants of this recovery. Assuming that this recovery is of a similar magnitude to the cumulative oil price decline since June 2014 and is composed of similar oil demand and oil supply shocks, all indications are that the U.S. economy's response would be largely symmetric. For example, one would expect a negative stimulus from consumer and non-oil investment spending. For the reasons discussed above, the positive stimulus to oil-related investment may be not quite as strong as the negative stimulus we saw during the 2014:Q3–2016:Q1 period; but even in that case, the net effect on the economy would be near zero. Of course, there is no reason to expect the composition, magnitude, and evolution of the oil demand and oil supply shocks to mirror those in the past. For example, if a recovery of the real price of oil primarily reflected a more robust global economy, the overall effects on the U.S. economy would be less negative than if the oil price recovery were driven mainly by actual or anticipated oil supply shocks.

ACKNOWLEDGMENTS We thank Richard Curtin at the University of Michigan's Institute for Social Research for providing access to additional data on gasoline price expectations in the Michigan Survey of Consumers that are not in the public domain. We also benefited from helpful comments by David Chasman, Lucas Davis, James Hamilton, Ana Maria Herrera, Valerie Ramey, and James Stock.

References

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. "What Do Consumers Believe about Future Gasoline Prices?" *Journal of Environmental Economics and Management* 66, no. 3: 383–403.

Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims. 2013. "Uncertainty and Economic Activity: Evidence from Business Survey Data." *American Economic Journal: Macroeconomics* 5, no. 2: 217–49.

Baumeister, Christiane, and Lutz Kilian. 2016a. "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us." *Journal of Economic Perspectives* 30, no. 1: 139–60.

———. 2016b. "Understanding the Decline in the Price of Oil since June 2014." *Journal of the Association of Environmental and Resource Economists* 3, no. 1: 131–58.

Baumeister, Christiane, Lutz Kilian, and Thomas K. Lee. 2016. "Inside the Crystal Ball: New Approaches to Predicting the Gasoline Price at the Pump." Forthcoming, *Journal of Applied Econometrics.*

Bernanke, Ben S. 1983. "Irreversibility, Uncertainty, and Cyclical Investment." *Quarterly Journal of Economics* 98, no. 1: 85–106.

Bresnahan, Timothy F., and Valerie A. Ramey. 1993. "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry." *American Economic Review* 83, no. 2: 213–18.

Coglianese, John, Lucas W. Davis, Lutz Kilian, and James H. Stock. 2016. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." Forthcoming, *Journal of Applied Econometrics.*

Davis, Steven J., and John Haltiwanger. 2001. "Sectoral Job Creation and Destruction Responses to Oil Price Changes." *Journal of Monetary Economics* 48, no. 3: 465–512.

Edelstein, Paul, and Lutz Kilian. 2007. "The Response of Business Fixed Investment to Changes in Energy Prices: A Test of Some Hypotheses about the Transmission of Energy Price Shocks." *B.E. Journal of Macroeconomics* 7, no. 1: article 35.

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

Elder, John, and Apostolos Serletis. 2010. "Oil Price Uncertainty." *Journal of Money, Credit and Banking* 42, no. 6: 1137–59.

Feyrer, James, Erin T. Mansur, and Bruce Sacerdote. 2015. "Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution." Working Paper no. 21624. Cambridge, Mass.: National Bureau of Economic Research.

Hamilton, James D. 1988. "A Neoclassical Model of Unemployment and the Business Cycle." *Journal of Political Economy* 96, no. 3: 593–617.

———. 2009. "Causes and Consequences of the Oil Shock of 2007–08." *Brookings Papers on Economic Activity,* Spring: 215–61.

———. 2013. "Oil Prices, Exhaustible Resources and Economic Growth." In *Handbook on Energy and Climate Change,* edited by Roger Fouquet. Cheltenham, U.K.: Edward Elgar.

Hamilton, James D., and Michael T. Owyang. 2012. "The Propagation of Regional Recessions." *Review of Economics and Statistics* 94, no. 4: 935–47.

Herrera, Ana María, and Mohamad B. Karaki. 2015. "The Effects of Oil Price Shocks on Job Reallocation." *Journal of Economic Dynamics and Control* 61: 95–113.

Herrera, Ana María, Mohamad B. Karaki, and Sandeep Kumar Rangaraju. 2016. "Where Do Jobs Go When Oil Prices Drop?" Forthcoming, *Energy Economics.*

Herrera, Ana María, Latika Gupta Lagalo, and Tatsuma Wada. 2011. "Oil Price Shocks and Industrial Production: Is the Relationship Linear?" *Macroeconomic Dynamics* 15, suppl. 3: 472–97.

———. 2015. "Asymmetries in the Response of Economic Activity to Oil Price Increases and Decreases?" *Journal of International Money and Finance* 50: 108–33.

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.

Kellogg, Ryan. 2014. "The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling." *American Economic Review* 104, no. 6: 1698–734.

Kilian, Lutz. 2009. "Comment on 'Causes and Consequences of the Oil Shock of 2007–08.'" *Brookings Papers on Economic Activity,* Spring: 267–78.

———. 2010. "Explaining Fluctuations in Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market." *Energy Journal* 31, no. 2: 87–112.

———. 2014. "Oil Price Shocks: Causes and Consequences." *Annual Review of Resource Economics* 6: 133–54.

———. 2016a. "The Impact of the Fracking Boom on Arab Oil Producers." Forthcoming, *Energy Journal*.

———. 2016b. "The Impact of the Shale Oil Revolution on U.S. Oil and Gasoline Prices." *Review of Environmental Economics and Policy* 10, no. 2: 185–205.

Kilian, Lutz, and Cheolbeom Park. 2009. "The Impact of Oil Price Shocks on the U.S. Stock Market." *International Economic Review* 50, no. 4: 1267–87.

Kilian, Lutz, and Clara Vega. 2011. "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices." *Review of Economics and Statistics* 93, no. 2: 660–71.

Kilian, Lutz, and Robert J. Vigfusson. 2016. "The Role of Oil Price Shocks in Causing U.S. Recessions." Forthcoming, *Journal of Money, Credit and Banking*.

Kleinberg, Robert L., Sergey Paltsev, Charles K. Ebinger, David Hobbs, and Tim Boersma. 2016. "Tight Oil Development Economics: Benchmarks, Breakeven Points, and Inelasticities." Working Paper no. 2016-012. Cambridge, Mass.: Massachusetts Institute of Technology, Center for Energy and Environmental Policy Research.

Lee, Kiseok, and Shawn Ni. 2002. "On the Dynamic Effects of Oil Price Shocks: A Study Using Industry-Level Data." *Journal of Monetary Economics* 49, no. 4: 823–52.

Pindyck, Robert S. 1991. "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature* 29, no. 3: 1110–48.

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.

Venditti, Fabrizio. 2013. "From Oil to Consumer Energy Prices: How Much Asymmetry along the Way?" *Energy Economics* 40: 468–73.

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.

Zarnowitz, Victor, and Louis A. Lambros. 1987. "Consensus and Uncertainty in Economic Prediction." *Journal of Political Economy* 95, no. 3: 591–621.