The Relationship Between Fossil Fuel Energy Consumption and Long-Run Growth: A Vector Autoregressive Approach

A Senior Comprehensive Paper

presented to
the Faculty of Carleton College
Department of Economics
Professor Stephie Fried, Advisor

by
Il Shan Ng

February 24, 2017

Abstract

I use structural vector autoregression to investigate how a permanent reduction in fossil fuel energy consumption affects long-run growth in the U.S. I first model the dynamic relationships among capital, fossil fuel energy consumption and total factor productivity. I then simulate the effects of a negative shock to fossil fuel energy consumption to understand the channels through which the shock may affect output. Finally, I derive the impulse response for output to examine its post-shock behavior. I find that although the growth rate falls in the short-run, it quickly returns to trend within 1.5 years of the shock, leaving long-run growth unaffected. Moreover, the short-run output gap is driven by productivity declines and not by contractions in the capital stock. My results suggest that the temporary economic losses induced by climate policy can be mitigated by improving productivity and more importantly, that climate policy may be pursued without compromising long-run growth.

Contents

1. Introduction

1.1 Research Question and Motivation

Much of the opposition against climate policy roots itself in the belief that cutting emissions to combat climate change necessarily entails slower economic growth and stifled development. In the 2016 Gallup poll about the environment, 37 percent of respondents indicated that economic growth should be given priority over environmental protection.[1] That is, over a third of Americans believe that climate policy and growth cannot be pursued hand in hand. While this may be true of the short-run, the long-run impact of such policy on growth is less clear. In this paper, I use vector autoregression to investigate the long-run behavior of growth in response to a permanent reduction to fossil fuel energy consumption in the U.S. I hypothesize that long-run growth is unaffected by this negative shock, and a confirmation of my hypothesis would constitute evidence that climate policy brings about economic losses that are only temporary.

The motivation for my paper is summarized in Figure 1.1, which shows how U.S. fossil fuel energy consumption and real GDP have evolved over the post-war period 1949 to 2015.[2] Both time series generally trend upward, giving the visual impression that output is positively correlated with the amount of energy consumed. There are two major exceptions to this trend, however, in the two time periods 1972–85 and 2007–15, during which fossil fuel energy consumption fell but output continued to rise. In fact, the general increase in output as a

[1]Taken from http://www.gallup.com/poll/1615/environment.aspx, last accessed on February 22, 2017.
[2]I obtained data on both variables from the Federal Reserve Bank of St. Louis website.

whole seems fairly invariant even to drastic declines in energy consumption. This led me to suspect that while negative shocks to fossil fuel energy consumption may induce lower output through short-run business cycle effects, they exert no significant impact on the long-run growth rate of output.

Figure 1.1: Fossil fuel energy consumption and output in the U.S.
(exceptions to the trend are highlighted in gray)

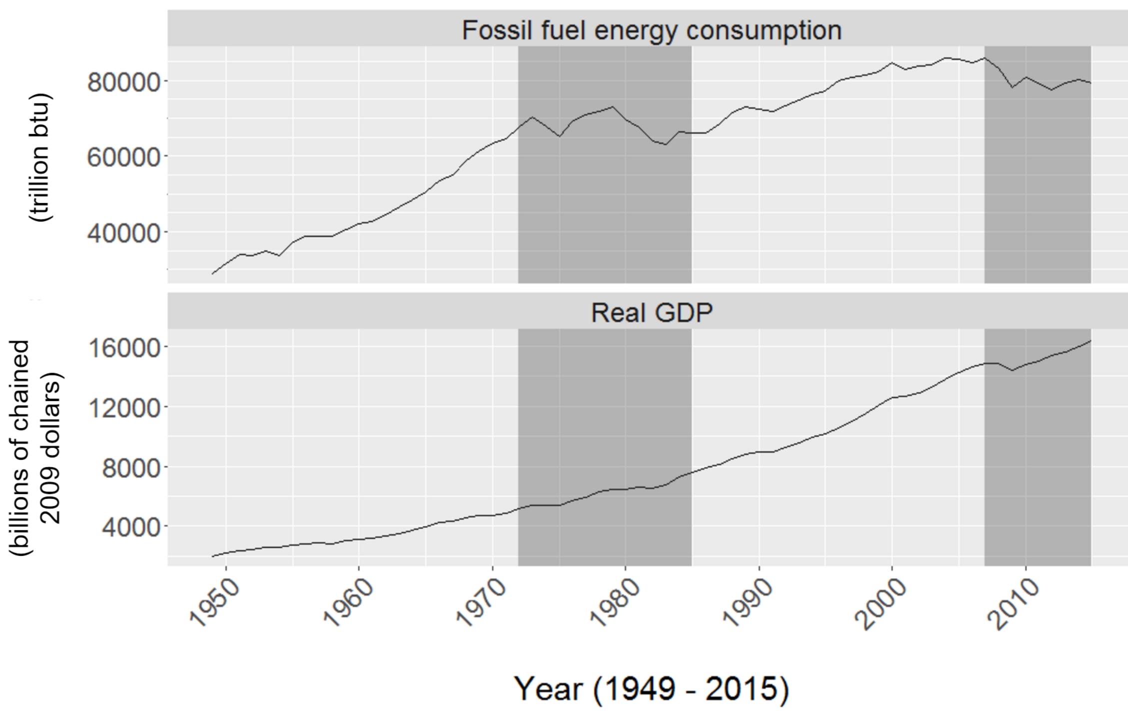

The literature on the time-dependent relationship between energy consumption and output, the so-called output-energy nexus, has not progressed beyond testing for causality between the time series of these two variables. In my paper, I move beyond this discussion of causality and focus on the structural effects of a sudden decrease in fossil fuel energy consumption. I first decompose output into its factor inputs and create a fully-identified vector autoregressive model that captures the dynamic relationships among them. I then simulate the effects of a negative exogenous shock to fossil fuel energy consumption on the other factor inputs, and use the impulse response for each factor input to understand the channels through which this shock may affect output. Finally, I derive the impulse response for output from the impulse responses for the individual factor inputs to determine the overall effect of the shock on output.

From the reduced form estimates and impulse responses for the individual factor inputs, I find that the negative shock to fossil fuel energy consumption causes total factor productivity to fall temporarily but exerts no significant effects on capital. The derived impulse response for output shows that the growth rate of output falls in the short-run but that this decrease is statistically significant for only about one year after the shock. Together, these observations imply that the economic losses brought about by climate policy are driven more by temporary productivity declines and not contractions in the size of the capital stock, and that long-run growth remains unaffected. Any fear that climate policy irreversibly slows down economic growth is therefore largely unfounded.

1.2 Vector Autoregression

The crux of my paper is the use of vector autoregression (VAR), a multivariate generalization of the autoregressive model that explains the evolution of one variable in terms of both its past values and the past values of other correlated variables. As my empirical analysis relies on theoretically-supported inter-causalities between the factor inputs, VAR is the appropriate econometric technique to use because it is able to model these bi-directional relationships as linear interdependencies. In my analysis, I begin by estimating the VAR model's reduced form, which assumes the absence of any contemporaneous effects of one endogenous variable on another and expresses each variable as a function of only pre-determined past values. Doing so yields consistent parameter estimates via ordinary least squares (OLS) estimation, but because all the error terms are correlated in the reduced form, it cannot be used to simulate the effect of an exogenous shock to any one variable.

Instead, I map my reduced form estimates onto the structural form, which allows variables to exert contemporaneous effects on each other and more importantly, features uncorrelated error terms. After proper identification to remove any simultaneity bias,[3] the orthogonalized impulse responses isolate the effects of the structural shock to fossil fuel energy consumption on the other factor inputs. This structural VAR (SVAR) model is therefore the preferred tool for policy analysis in general (Kim, Heyongwoo, 2012), and is what gives me my main results.

The remainder of the paper is organized as follows. I first present a literature review of previous work on the output-energy nexus, and explain what my comprehensive exercise contributes to the literature. I then delve into the economic theory that informs my empirical model and allows me to decompose output into its factor inputs. Next, I describe the data and any data transformations that I perform to render the data in a form suitable for time series analysis. Following this, I present my empirical results and perform robustness tests on my model. Finally, I conclude with any implications that my results may have for climate policy and propose avenues for future research.

[3] This is also known as an endogeneity problem, and occurs when an explanatory variable is correlated with the error term in the same equation.

2. Literature Review

What makes the output-energy nexus the subject of nearly four decades of intense research is the inherent inter-causality between the two variables. Because energy is an important factor input in the production process, higher energy consumption tends to drive economic growth. Likewise, a higher level of economic development may lead to more efficient energy use, which in turn influences energy consumption. As a result, the direction of causality between output and energy cannot be determined *a priori*, which leaves empirical work as the only other option (Halicioglu, 2008).

Kraft and Kraft (1978) pioneered this empirical work by applying Granger causality tests to data on the post war period 1947–1974 and finding that output Granger-causes energy consumption but not the other way around.[4] This inspired an explosion of further work on the output-energy nexus, with researchers applying different causality tests to different time periods in different countries. Among the techniques that were brought to bear on this question is VAR, which captured the dynamic relationship between output and energy models better than static models did. Unfortunately, much of the literature has not moved on from attempts to detect causality between output and energy consumption because of highly conflicting results. Based on the direction of causality found by the study, these differing results may be classified into one of four hypotheses which each has a different implication for energy policy (Ozturk 2010):

[4]Granger causality is a statistical concept of causality that is based on the ability of one signal to predict another signal. If a time series X Granger-causes another time series Y, then the lagged values of X contains information that can help predict Y above and beyond the information contained in the lagged values of Y alone.

I The neutrality hypothesis. There is no causality between output and energy consumption, which implies that energy policy will have no significant effect on economic growth.

II The conservation hypothesis. There is unidirectional causality running from output to energy consumption, suggesting that energy policy may pursue energy conservation with no significant adverse impacts on economic growth.

III The growth hypothesis. There is unidirectional causality running from energy consumption to output. This hypothesis perceives energy as an important complement to labor and capital in the production process. Hence, energy policies that attempt to limit energy consumption will tend to exert undesirable effects on economic growth.

IV The feedback hypothesis. There is bi-directional causality between output and energy consumption. This suggests a simultaneity problem that may require energy policy to adapt quickly to changes in either variable. For example, limiting energy consumption would reduce output which could in turn further decrease energy consumption.

These diverse results arise partly because of the choice of different countries among these studies; causality might run in different directions in different countries due to disparities in energy fuel mix, indigenous energy sources, institutional arrangements and energy policy, among other factors (Chen et al., 2007). However, even studies that focused on the same country have managed to produce wildly different results. Between 1978 and 2009, of the eight studies that focused on the U.S., five found no causality between output and energy consumption (Akarca and Long, 1980; Yu and Hwang, 1984; Yu and Jin, 1992; Cheng, 1995; Payne, 2009), two found unidirectional causality running from output to energy consumption

(Kraft and Kraft, 1978; Abosedra and Baghestani, 1989), and one found unidirectional causality running from energy consumption to output (Bowden and Payne, 2009).

Although a thorough comparison of the econometric models used and battery of statistical tests employed by each of these studies is beyond the scope of this literature review,[5] I am able to distill three important factors that make the results from a study more reliable. First, as Ozturk (2010) asserts, VAR models that capture the dynamic interdependencies among the modeled variables tended to perform better than static models that rely simply on cross-sectional analyses. Second, multivariate models that include not only output and energy consumption but also other factors of production tended to provide a better representation of real-world interactions by accommodating several mechanisms and causality channels. Ghali and El-Sakka (2004), for instance, modeled a four-variable system involving output, capital, labor and energy consumption in Canada, and concluded that energy is an important limiting factor for economic growth in Canada. Lastly, as applies to time series analysis in general, studies that used more time points over a longer time period tended to produce more reliable results.

I incorporate these three observations into my analysis. The core of my research is my VAR model, and instead of modeling a bivariate system of just output and energy consumption, I decompose output into its factors of production and attempt to model any causalities between energy consumption and these other factor inputs. Doing so not only reduces any potential variable omission bias that plagued previous bivariate studies, but also gives me insight into the specific channels through which changes in energy consumption may affect output. To

[5] Some of the techniques used by authors studying output and energy consumption in the U.S. are Vector Error Correction Methods, Johansen cointegration and the Toda-Yamamoto test for causality, all of which are fairly advanced and require a graduate-level understanding of time series econometrics.

ensure that my sample size exceeds those of previous studies, I use all publicly available data to include the years that come after these previous studies.

Because I model the relationships between energy consumption and the other factor inputs and not output, my contribution to the literature will not be another test for the direction of causality between energy consumption and output. Rather, I sidestep the causality debate and focus on developing the structural component of my model. Doing so allows me to simulate the economic impacts of a negative shock to energy consumption and predict the shock's effects on the other factor inputs and therefore output. This structural analysis is more appropriate for policy analysis because a mandated cut in energy use is better thought of as an exogenous, unanticipated shock to energy consumption instead of an endogenous change (Ghali and El-Sakka, 2004). The orthogonalized impulse responses from a SVAR model also describe real-world economic channels following a shock more accurately than the generalized impulse responses from reduced form models, which are invariant to how the endogenous variables are ordered (Kim, Hyeongwoo, 2012). And so, by moving past the causality issue and into structural analysis, my paper is in fact among the first to design a fully-specified model that can be used to assess energy policy.

3. Economic Theory

3.1 Growth Rate Decomposition

In this section, I establish the relationship between fossil fuel energy consumption and real output. Because fossil fuel energy is a factor input in the production process, a useful starting point is the theory of production. I follow the standard in the literature on climate-economy systems (Hassler, Krusell and Olovsson, 2012; Werf, 2008) and use a Cobb-Douglas production function with constant returns to scale to decompose output into its factor inputs. Let Y_t represent U.S. real GDP in time t and write it as

$$Y_t = A_t K_t^{\beta} E_t^{\gamma} L_t^{1-\beta-\gamma}, \tag{3.1}$$

where A_t is total factor productivity (TFP), K_t is the capital stock, E_t is fossil fuel energy consumption, and L_t is the size of the labor force all in year t. Assuming that all factor markets are perfectly competitive, the factor shares of income β, γ and $1-\beta-\gamma$ are constant (Hassler, Krusell and Olovsson, 2012). I set the value of β, the capital share of income, to 0.33, a number empirically determined by the literature and which macroeconomists generally accept to be true for the U.S. For the value of γ, the energy share of income in the U.S., I use the average estimate of 0.04 from Golosov, Hassler and Krusell (2012).[6]

I now attempt to express the growth rate of output in terms of the growth rate of its

[6]Golosov et al. actually accounted for both fossil fuel and non-fossil-fuel energy sources in producing the estimate of 0.04. However, because fossil fuels constitute the overwhelming majority (91 percent) of energy consumption, 0.04 is still a good approximation when only considering fossil fuel energy.

factor inputs. Applying the natural logarithm to both sides of Equation 3.1 gives

$$\ln Y_t = \ln A_t + \beta \ln K_t + \gamma \ln E_t + (1 - \beta - \gamma) \ln L_t,$$

and taking the derivative of both sides with respect to t yields

$$\frac{1}{Y_t}\frac{dY_t}{dt} = \frac{1}{A_t}\frac{dA_t}{dt} + \beta\frac{1}{K_t}\frac{dK_t}{dt} + \gamma\frac{1}{E_t}\frac{dE_t}{dt} + (1 - \beta - \gamma)\frac{1}{L_t}\frac{dL_t}{dt}. \tag{3.2}$$

Because I define time in terms of discrete years, the differential dX_t/dt for any factor input X_t can be approximated by its year-on-year change ΔX_{t+1}, and so Equation 3.2 is approximately

$$\frac{\Delta Y_{t+1}}{Y_t} = \frac{\Delta A_{t+1}}{A_t} + \beta\frac{\Delta K_{t+1}}{K_t} + \gamma\frac{\Delta E_{t+1}}{E_t} + (1 - \beta - \gamma)\frac{\Delta L_{t+1}}{L_t}. \tag{3.3}$$

Notice that $\Delta X_{t+1}/X_t$ is just the growth rate $g_{X,t}$ of the factor input X in year t. We can then re-write Equation 3.3 as

$$g_{Y,t} = g_{A,t} + \beta g_{K,t} + \gamma g_{E,t} + (1 - \beta - \gamma) g_{L,t}. \tag{3.4}$$

This suggests a multivariate analysis involving four variables. Unfortunately, performing VAR with four variables leads to a parameter identification problem whose solution is beyond the scope of this comprehensive exercise, and so I make a simplifying assumption here to fix the growth rate in one of these variables. In particular, because I am interested in the long-run behavior of the system and not on short-run business cycles, I assume the growth rate of labor to be constant. I can therefore remove labor from my empirical model and base my analysis on growth rates of only the three variables TFP, capital and fossil fuel energy consumption.

Let $\Delta g_{Y,t}, \Delta g_{A,t}, \Delta g_{K,t}$ and $\Delta g_{E,t}$ represent deviations in the growth rates of the respective variable t years after a one-time negative shock to the growth rate of fossil fuel energy consumption (that is, the difference between the growth rates for a world without the shock, and a world with a shock). Then, we can re-express Equation 3.4 as the linear combination

$$\Delta g_{Y,t} = \Delta g_{A,t} + \beta \Delta g_{K,t} + \gamma \Delta g_{E,t}. \tag{3.5}$$

The main goal of my VAR analysis is therefore to estimate $\Delta g_{A,t}, \Delta g_{K,t}$ and $\Delta g_{E,t}$ in the long-run and sum them using appropriate weights according to Equation 3.5 to derive $\Delta g_{Y,t}$. Finding that $\Delta g_{Y,t}$ significantly differs from zero would lead to a rejection of my hypothesis that a one-time reduction in the growth rate of fossil fuel energy consumption does not affect long-run growth.

3.2 Interdependencies Among Variables

A key assumption in the use of VAR is that the variables in the system may be correlated with each other across different time periods, such that the evolution of one variable over time can be explained partially by the lagged values of the other variables. In this section, I present relevant theory to argue for possible interdependencies among my three variables TFP, capital and fossil fuel energy consumption.

To see how changes in TFP may affect the amount of capital and fossil fuel energy used in an economy, I once again turn to the theory of production. Taking the first derivative of Equation 3.1 with respect to capital and fossil fuel energy gives the respective marginal products

$$MP_K = A_t \beta K_t^{\beta-1} E_t^{\gamma} L_t^{1-\beta-\gamma} \qquad \text{and} \qquad MP_E = A_t K_t^{\beta} \gamma E_t^{\gamma-1} L_t^{1-\beta-\gamma}.$$

Assuming the case of perfect competition, profit-maximizing firms hire capital up to the point where the real rate of return on capital equals the marginal product of capital. Similarly, firms consume fossil fuel energy up to the point where the price of fossil fuel energy equals the marginal product of fossil fuel energy. Letting r denote the real return on capital and P_E denote the price of fossil fuel energy, the following therefore hold at equilibrium:

$$r = A_t \beta K_t^{\beta-1} E_t^{\gamma} L_t^{1-\beta-\gamma} \qquad \text{and} \qquad P_E = A_t K_t^{\beta} \gamma E_t^{\gamma-1} L_t^{1-\beta-\gamma}.$$

Rearranging the above to express quantity of the input in terms of its price gives

$$K_t = r \left(A_t \beta E_t^{\gamma} L_t^{1-\beta-\gamma} \right)^{\frac{1}{1-\beta}} \qquad \text{and} \qquad E_t = P_E \left(A_t K_t^{\beta} \gamma L_t^{1-\beta-\gamma} \right)^{\frac{1}{1-\gamma}}, \tag{3.6}$$

which are the demand curves for capital and fossil fuel energy respectively. Notice that $\frac{1}{1-\beta} > 0$ and $\frac{1}{1-\gamma} > 0$, and so an increase in TFP would cause both of these demand curves to shift rightward. Assuming that the respective supply curves remain unchanged, the amount of capital hired and fossil fuel energy consumed at equilibrium should therefore increase.

To see how changes in the size of the capital stock might affect the amount of fossil fuel energy consumed, I use the idea that capital and energy are economic complements in the production process (Prywes, 1986). Because capital requires energy to operate, when firms acquire more capital, they naturally demand more energy. This is also reflected more precisely in Equation 3.6; an increase in capital would cause the demand curve for fossil fuel energy to shift rightward. If the supply curve for fossil fuel energy remains unchanged, firms will now consume more fossil fuel energy at equilibrium. Another school of thought, however, believes that capital and energy are economic substitutes. Koete, Groot and Florax (2008) conducted a meta-analysis of capital-energy substitutions and concluded that elasticity estimates from the literature tended to indicate substantial capital-energy substitutability, particularly in the long-run. The intuition is that given enough time, innovating firms will hire more energy-efficient capital to make up for the lower availability of energy. Sorting out these two competing views is beyond the scope of this paper; this disagreement, however, suggests that the correlation between capital and energy could be positive or negative, or possibly even zero.

Next, to explain how changes in the size of the capital stock may affect TFP, I rely on the notion of idle capital from the theory of capital utilization (Winston, 1974). According to this idea, firms may leave a fraction of their capital idle in times of deficient demand for their goods. Because TFP is an aggregate measure of how efficiently producers in an economy use their factor inputs, the inability of firms to use their capital at full capacity reduces TFP. When firms acquire more capital, their tendency to leave capital idle due to insufficient demand increases by virtue of the now larger capital stock. In other words, the more capital firms own, the harder it becomes for these firms to run their capital at full capacity all of the time. As a result, TFP decreases.

I can likewise use the theory of capital utilization to explain how changes in fossil fuel energy consumption may influence TFP. When firms are forced to reduce their consumption of fossil fuel energy — by climate policy, for example — they may have to shut down a fraction of their capital. Because the size of the capital stock tends to be fixed in the short-run, these firms may be unable to retire the idle capital quickly enough. This leads to excess capacity, which causes TFP to fall. Finally, as described earlier, the disparate views on whether energy and capital are complements or substitutes make it difficult to predict how changes in fossil fuel energy consumption might affect capital. For simplicity, I assume that these two competing mechanisms are equally present, such that there is zero correlation between fossil fuel energy and capital.

Table 3.1 summarizes the foregoing discussion on the interdependencies among my three variables. Because a secondary focus of my analysis is to examine how a negative shock to fossil fuel energy consumption affects the other factors of production, I make my final two conjectures sub-hypotheses which I will test with my VAR results. I use the other conjectures to assess whether my model produces results that are consistent with economic theory — this will take the place of VAR-specific diagnostic tests which tend to be less contextual and harder to implement.[7]

Table 3.1: Summary of interdependencies among variables

		Effect variable		
		TFP	Capital	Energy
Cause variable	TFP	-	Increase in TFP increases capital	Increase in TFP increases energy
	Capital	Increase in capital decreases TFP	-	Zero correlation
	Energy	Decrease in energy decreases TFP (sub-hypothesis 1)	Zero correlation (sub-hypothesis 2)	-

[7]Some diagnostic tests for VAR models include forecast error variance decompositions, checking for stability by examining the eigenvalues of the dynamic matrix, and testing for Granger causality. Of these, testing for Granger causality is the most straightforward, and I do so implicitly by checking the signs and magnitudes of the estimated coefficients of the reduced form VAR.

4. Data and Methodology

4.1 Exploratory Data Analysis

I obtained time series data on fossil fuel energy consumption from the U.S. Energy Information Administration, and time series data on capital in the form of real gross private domestic investment from the Federal Reserve Bank of St. Louis (FRED) website. Because data on TFP was not readily available, I derived the time series for TFP using the Cobb-Douglas specification in Equation 3.1 and data on U.S. real GDP and size of the labor force, both of which I had also obtained from FRED. Figure 4.1 shows plots of the time series data for the three variables over the common time period 1949 – 2015.

Figure 4.1: Time series plots for (a) TFP, (b) capital and (c) fossil fuel energy consumption

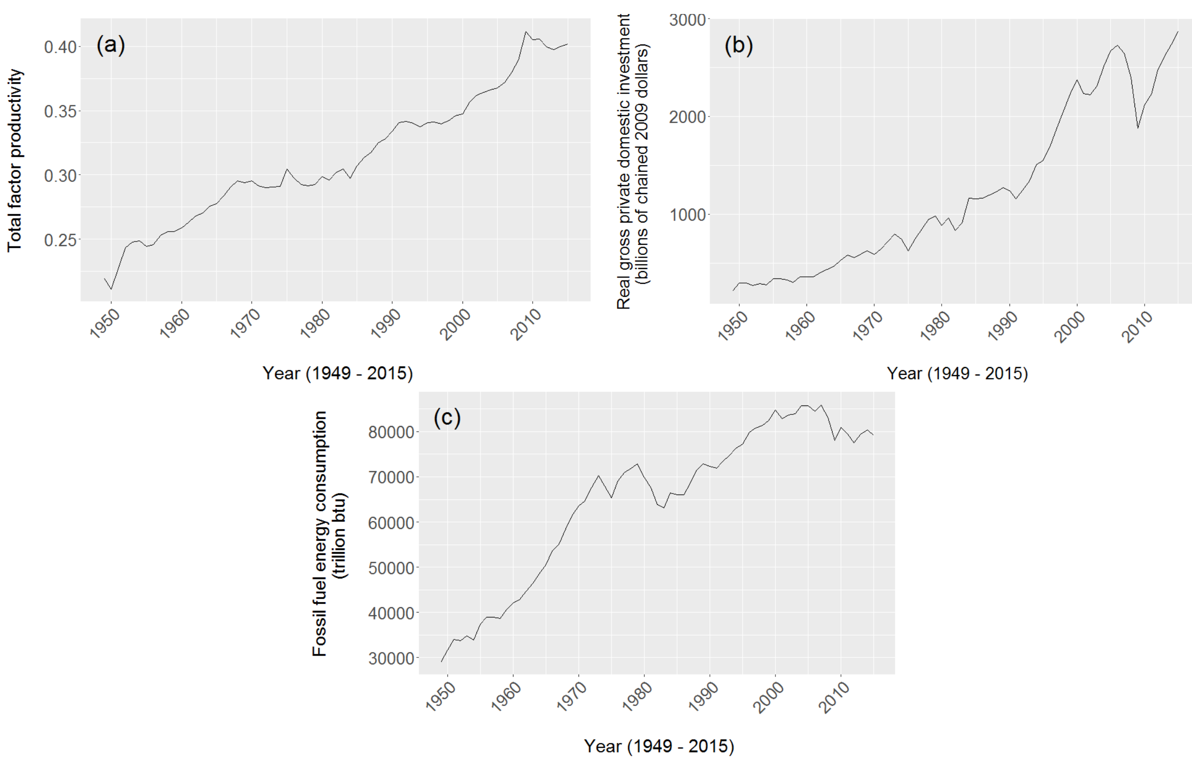

Because I intend to work with the growth rates of each of the variables, I apply the natural logarithm transformation to each of the time series. Figure 4.2 shows the transformed data. There is sufficient variation (in the form of fluctuations about the trend line) for each logged variable, such that the VAR is likely to detect significant coefficient estimates. More importantly, these plots show that all three logged variables have non-stationary, trending series. Because the OLS estimation procedures in time series analysis rely on the data being stationary, I perform a series of preliminary tests in the next section to figure out how best to further transform the data to achieve stationarity for each variable.

Figure 4.2: Time series plots for (a) log of TFP, (b) log of capital and (c) log of fossil fuel energy consumption

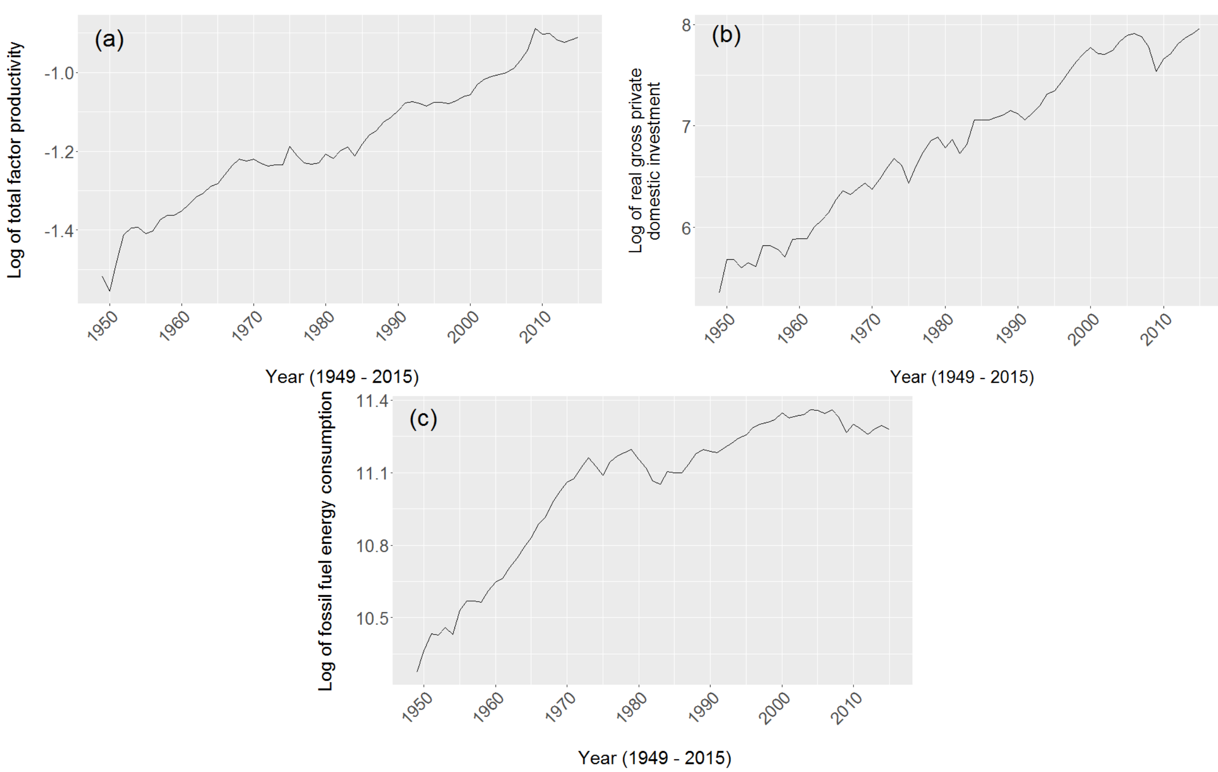

4.2 Preliminary Tests

There are two main ways to achieve stationarity for time series data: de-trending and first-differencing. Which of these two methods to choose depends on whether the non-stationary time series concerned follows a deterministic or stochastic trend. To answer this question, I perform the Dickey-Fuller test for the presence of a unit root on the time series for each variable separately. The null hypothesis is that the univariate time series has a unit root and therefore follows a stochastic trend. The alternative hypothesis is that the univariate time series has no unit root and is therefore trend-stationary (i.e. follows a deterministic trend). Table 4.1 shows results of the test performed in the statistical software package R. By comparing the computed test statistics to the critical values for the Dickey-Fuller distribution, I reject the null hypothesis for the log of TFP at the 1 percent level of significance. The time series for the log of TFP is trend-stationary, whereas those for the log of capital and fossil fuel energy exhibit stochastic trends.

Table 4.1: Results for the Dickey-Fuller test for unit roots

	log of TFP	log of capital	log of energy
Test statistic	-5.014	-3.051	-1.169

The critical values are: -4.04 (1 percent), -3.45 (5 percent), -3.15 (10 percent).

Because the series for the log of TFP is trend-stationary, I can make it stationary simply by removing the trend. I do so by performing a simple OLS regression of the log of TFP against *year* and recovering the residuals. Because the *year* variable controls for the upward deterministic trend, the residuals should ideally oscillate with constant variance about a mean of zero. Figure 4.3 shows a plot of these residuals against time. With the exception of the year 1950, the transformed time series does look stationary. To make this observation more formal, I perform another Dickey-Fuller test on this de-trended series, except this time I use the alternative hypothesis that the series is stationary (as opposed to trend-stationary). The computed test statistic is -3.36, which does not exceed the 1 percent critical value of -2.6, allowing me to reject the null hypothesis for the presence of a unit root and conclude that this de-trended series is indeed stationary.

Figure 4.3: Time series for de-trended log of TFP

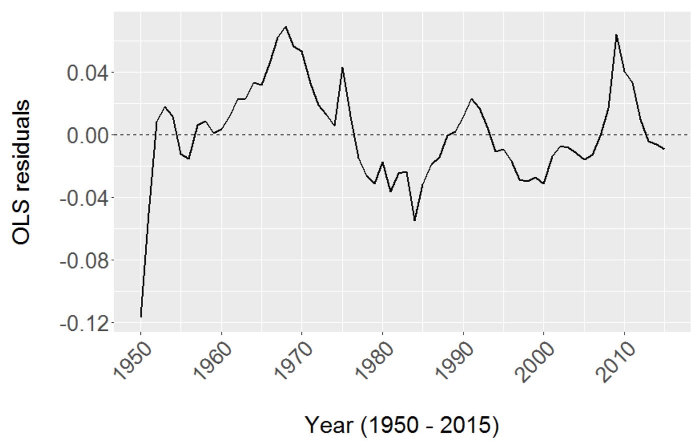

For the other two time series, which have stochastic trends, first-differencing is necessary to achieve stationarity in each of the variables. However, when modeling two or more time series with stochastic trends in a multivariate framework, one has to worry about whether or not the time series variables are co-integrated. Formally, a collection of time series variables are co-integrated if each of them can be made stationary via first-differencing,[8] but there exists a linear combination of the variables that is stationary. Co-integrated variables should always be modeled in levels to avoid spurious regressions.[9]

[8]This is more commonly known in the literature as being integrated of order one.

[9]The basic idea is this: even though the co-integrated variables are individually non-stationary, OLS regression will produce coefficient estimates such that the stochastic trends cancel each other out. These variables therefore become stationary when modeled together.

To determine whether the time series for the log of capital and fossil fuel energy are co-integrated, I apply the Phillips-Ouliaris residual-based test for co-integration.[10] The test first performs OLS estimation of the co-integrating equation[11] and recovers the residuals. Under the null hypothesis that the two time series are not co-integrated, *all* linear combinations of the two variables, including the estimated residuals, are non-stationary. Therefore, a test of the null hypothesis of no co-integration amounts to a unit root test on the residuals of the co-integrating equation (Phillips and Ouliaris, 1990). Using R, I found the p-value for the Phillips-Ouliaris test to be bounded below by 0.15, and so I do not reject the null hypothesis that the two variables are not co-integrated.

First-differencing is therefore the appropriate transformation for the remaining two time series. For a generic time series variable X_t, define its first-difference as

$$\nabla X_t = X_t - X_{t-1}.$$

If X_t is a logged variable, then its first-difference approximates its year-on-year growth rate — a nice economic interpretation on top of achieving stationarity. Figure 4.4 shows plots of the first-differenced variables. Although the time series for capital now looks stationary, the series for fossil fuel energy shows some hint of a downward trend. To be sure, I run the Dickey-Fuller test on the first-differenced variables using the alternative hypothesis that the series is stationary and get a test statistic of -5.576 for capital and -4.622 for fossil fuel energy. Comparing these to the 1 percent critical value of -2.6 for the Dickey-Fuller distribution, I

[10] There are other more general tests for co-integration, such as the Johansen test which uses the trace or maximum eigenvalue of the coefficient matrix as the test statistic. Understanding these more complicated tests is beyond the scope of my comprehensive exercise, and I use the simpler Phillips-Ouliaris test.

[11] The co-integrating equation is the linear combination of the two variables that produces a stationary time series.

am able to reject the null hypothesis for both variables and conclude that first-differencing has indeed made both variables stationary.

Figure 4.4: Time series plots for the first difference of the log of (a) capital and (b) fossil fuel energy consumption

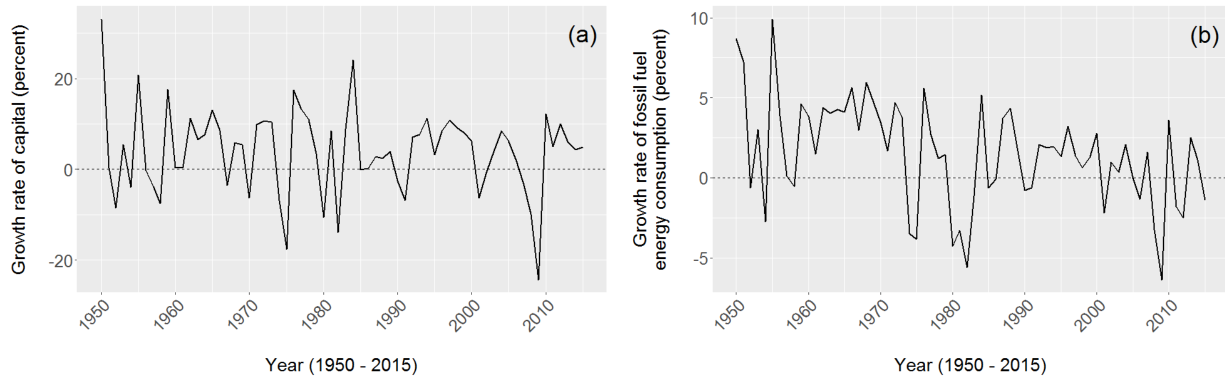

The last preliminary test involves finding the optimal lag length for my VAR model. Because the empirical literature is inconsistent in the number of lags used, and because I could not find any theoretical reasons to support the use of a certain number of lags,[12] I use Information Criteria to select the optimal number of lags for my data. Table 4.2 shows four information criteria computed for lag lengths one to five.[13] Recall that smaller Information Criteria values indicate better overall fit of the model to the data. Although the results for this test are split between one and three for the optimal lag length, I focus on the Bayes Information Criteria (BIC) because it is the most suitable for my model, which defines each time period as a year (Ivanov and Kilian, 2005). The BIC suggests an optimal lag length of one, which fortunately also helps me keep my model simple.

[12]For instance, if I had used quarterly data, then theory suggests that I should include at least four lags to capture the seasonal variation in energy consumption over one year.

[13]These four criteria are the Akaike Information Criteria (AIC), Hannan-Quinn Information Criteria (HQC), Bayes Information Criteria (BIC), and the Final Prediction Error (FPE).

Table 4.2: Information Criteria values for different lag lengths

	1	2	3	4	5
AIC	-22.13	-22.06	**-22.22**	-22.11	-22.21
HQC	**-21.92**	-21.74	-21.76	-21.54	-21.52
BIC	**-21.61**	-21.23	-21.07	-20.66	-20.45
FPE	2.46×10^{-10}	2.63×10^{-10}	$\mathbf{2.28 \times 10^{-10}}$	2.55×10^{-10}	2.36×10^{-10}

Bold values are the smallest value for each criterion, and correspond to the optimal lag length.

4.3 Empirical Model

Having attained stationarity in all my variables and determined the optimal lag length, I now specify my SVAR model. Let $\tilde{A}_t$ be the de-trended natural logarithm of TFP, and ∇K_t and ∇E_t be the first-differences of the natural logarithms of capital and fossil fuel energy consumption respectively. With a lag order of one year, the system of three equations to be estimated using OLS regression is

$$\tilde{A}_t = a_0 + a_1 \nabla K_t + a_2 \nabla E_t + a_3 \tilde{A}_{t-1} + a_4 \nabla K_{t-1} + a_5 \nabla E_{t-1} + \epsilon_t^A$$

$$\nabla K_t = b_0 + b_1 \tilde{A}_t + b_2 \nabla E_t + b_3 \tilde{A}_{t-1} + b_4 \nabla K_{t-1} + b_5 \nabla E_{t-1} + \epsilon_t^K$$

$$\nabla E_t = c_0 + c_1 \tilde{A}_t + c_2 \nabla K_t + c_3 \tilde{A}_{t-1} + c_4 \nabla K_{t-1} + c_5 \nabla E_{t-1} + \epsilon_t^E,$$

where $\epsilon_t^A, \epsilon_t^K$ and ϵ_t^E are the respective error terms with zero mean, zero pairwise covariance and zero correlation over time. As written, each error term is correlated with the regressors for the corresponding variable, presenting a simultaneity bias problem. For example, increasing ϵ_t^A would cause $\tilde{A}_t$ to increase, which in turn causes ∇K_t and ∇E_t to increase. But because ∇K_t and ∇E_t are covariates in the equation for $\tilde{A}_t$, they confound the effects of ϵ_t^A, such that the original change in ϵ_t^A cannot yet be interpreted as an exogenous shock to $\tilde{A}_t$. I resolve this problem by imposing zero restrictions on three coefficients in the above system using a few assumptions about the contemporaneous inter-dependencies among the three variables.

The first two restrictions are easy to obtain. Because the size of the capital stock in an economy is fixed in the short run, I assume that changes in TFP and fossil fuel energy consumption exert no contemporaneous effects on capital. The implies zero restrictions on b_1 and b_2, ordering capital first.[14] For the third restriction, I use the idea that firms may need some time to change their energy consumption patterns in response to a productivity shock, especially if they purchase energy via medium to long-term contracts or if changes in productivity are not immediately observable. This implies a zero restriction on c_1, ordering fossil fuel energy second and TFP third in the VAR.

The above restrictions ensure that each error term is no longer correlated with the regressors for the corresponding variable. Hence, a change in any of the error terms can be interpreted as an exogenous shock to the corresponding variable, and the model is now properly identified. In the context of my hypothesis, a permanent reduction in fossil fuel energy consumption can be interpreted as a one-time negative shock to ∇E_t, which approximates the growth rate of fossil fuel energy consumption. By graphing the impulse responses for all three variables following the shock, I can examine the dynamic effects of the shock on all three variables over time. This is the subject of the next chapter.

[14] A variable X is ordered before Y if changes in Y do not contemporaneously affect X; that is, if the regression coefficient of Y_t in the equation for X_t is zero.

5. Results and Discussion

5.1 Reduced Form Results

Table 5.1 presents the OLS regression results for the reduced form of my VAR model. The low multiple R-squared values for ∇K_t and ∇E_t immediately stand out as an issue and indicate that the endogenous variables in the system do not explain a large fraction of the observed variation in the growth rates of capital and fossil fuel energy consumption. Nonetheless, I check if the VAR is correctly specified by inspecting whether the sign and significance of each coefficient estimate matches up with the theoretical predictions discussed in Section 3.2.

Table 5.1: OLS regression results for the reduced form

Dependent Variable	$\tilde{A}_t$	∇K_t	∇E_t
Constant	0.000876 (0.00208)	0.0339*** (0.0121)	0.00931** (0.00423)
$\tilde{A}_{t-1}$	0.634*** (0.0673)	0.808** (0.391)	0.258* (0.137)
∇K_{t-1}	−0.0932*** (0.0313)	0.274 (0.182)	0.0266 (0.0634)
∇E_{t-1}	0.314*** (0.0847)	−0.683 (0.492)	0.219 (0.172)
N	65	65	65
Multiple R-squared	0.723	0.0750	0.124

Significance codes: * 10%, ** 5% *** 1%. Standard errors are given in parenthesis.

According to Table 5.1, TFP in the previous year is significantly and positively correlated with the growth rates of capital and fossil fuel energy consumption in the current year. This matches the summary given in the first row of Table 3.1. Furthermore, an increase in the previous year's growth rate of capital is associated with a decrease in TFP in the current year, but is uncorrelated with the growth rate of fossil fuel energy consumption in the current year. Again, this matches the summary given in the second row of Table 3.1. Therefore, even though the model suffers from low R-squared values, it at least produces coefficient estimates that align well with economic theory.

I now test my two sub-hypotheses first mentioned in the third row of Table 3.1. The first sub-hypothesis is that a decrease in fossil fuel energy consumption is associated with a decrease in TFP, because the former lowers capital utilization which in turn negatively affects the latter. Let x represent the true coefficient of ∇E_{t-1} in the reduced form equation for $\tilde{A}_t$. Then, the null and alternative hypotheses for this test are

$$H_0 : x = 0 \qquad \text{and} \qquad H_A : x > 0$$

respectively. Assuming that the null hypothesis is true, the test statistic $T = \frac{\hat{x}}{SE[\hat{x}]}$ follows a Student's t-distribution with 61 degrees of freedom (65 observations minus 4 parameter estimates). The observed test statistic is $t = \frac{0.314}{0.0847} = 3.707$, which, when compared to a $t_{df=61}$ distribution, yields a p-value of 0.000227. This allows me to reject the null hypothesis at the 1 percent level of significance, and conclude that there is strong evidence to suggest that a decrease in fossil fuel energy is indeed associated with a decrease in TFP.

The second sub-hypothesis is that fossil fuel energy consumption and capital are uncorrelated over different time periods, due to the competing mechanisms that treat energy and capital as substitutes or complements. Likewise, let y represent the true coefficient of ∇E_{t-1} in the reduced form equation for ∇K_{t-1}, and perform the following two-tailed test

$$H_0 : y = 0 \qquad \text{and} \qquad H_A : y \neq 0.$$

Under the null hypothesis, the test statistic $T = \frac{\hat{y}}{SE[\hat{y}]}$ follows a t-distribution with 61 degrees of freedom. The observed test statistic $t = \frac{-0.683}{0.492} = -1.388$ produces a two-sided p-value of 0.17, and so I do not reject the null hypothesis of zero correlation between fossil fuel energy consumption and capital. There is insufficient evidence to conclude that one of the two competing mechanisms outweigh the other, and so we still do not know whether energy and capital behave as substitutes or complements in the production process.

5.2 Impulse Responses

Figure 5.1 shows the impulse responses for 15 years following a 10 percent negative shock to the growth rate of fossil fuel energy consumption. In the top panel, the growth rate of fossil fuel energy consumption falls initially by 10 percent due to the exogenous shock, but recovers back to trend over the next 15 years. The narrowing 95 percent confidence intervals indicate that in almost all replicates, the deviation caused by the shock diminishes to zero.

The middle panel shows that TFP decreases below trend due to the negative shock to energy, and that this decrease is significant between years 0.5 and 5.5 at the 95 percent significance level, as indicated by the confidence bands. This provides further evidence to support my first sub-hypothesis that a negative shock to fossil fuel energy is associated with a fall in TFP. Like energy consumption, TFP recovers to trend within 15 years with narrowing confidence intervals. In fact, the decrease in TFP is no longer statistically significant 5.5 years after the shock, and so the productivity decline following a negative shock to energy consumption is fairly short-lived.

Figure 5.1: Impulse responses for a 10 percent negative shock to the growth rate of fossil fuel energy consumption (dashed lines give the bootstrapped 95 percent confidence intervals)

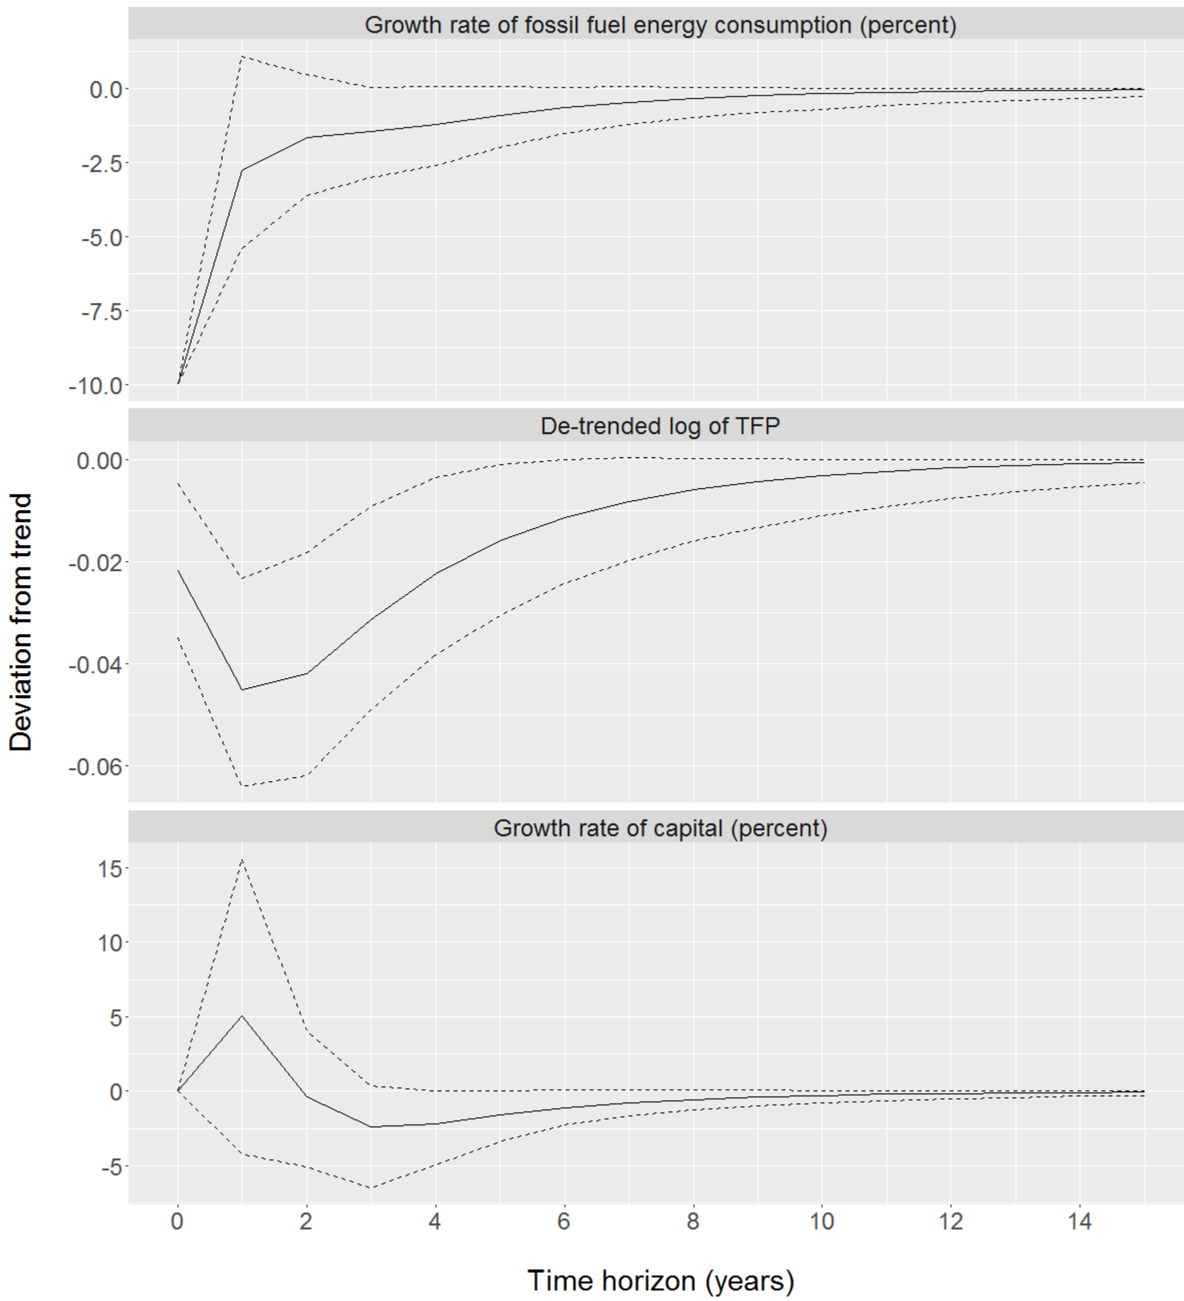

In the bottom panel, the growth rate of capital does not change immediately in the same time period as the energy shock. This reflects my identifying assumption that changes to fossil fuel energy consumption exert no contemporaneous effects on capital. However, the energy shock fails to bring about any statistically significant changes to the growth rate of capital even in the subsequent years. This lack of a dynamic effect is further evidence for my second sub-hypothesis that fossil fuel energy consumption and capital are uncorrelated across different time periods due to transmission channels that potentially conflict with each other.

The impulse responses for the individual variables provide visual evidence for my two sub-hypotheses, but do not directly answer my main question of how output behaves in the long-run following a negative shock to energy. Intuitively, because the deviations in all three factor inputs taper off to zero in Figure 5.1, output should return to trend within 15 years of the shock. To demonstrate this more formally, I derive the impulse response for output from the impulse responses for the individual factor inputs. Because I modeled energy and capital in terms of their growth rates but TFP in log-levels, this derivation is computationally tricky and I break it down into separate steps to ensure accuracy. First, I convert the impulse response for TFP into growth rates, combine this with the other two impulse responses, and use Equation 3.4 to derive the the growth rate of output in the world with the shock. I then use the reduced form estimates in Table 5.1 to forecast "business as usual" (BAU) growth rates for the factor inputs in a world without a shock to derive the BAU growth rate for output, again with Equation 3.4. For the growth rate of labor in Equation 3.4, I use the U.S. Bureau of Labor Statistics' projected rate of 0.5 percent over the 2014-24 period.[15]

Figure 5.2 plots the derived growth rates for output. In time period zero, the growth rate dips immediately to about -0.35 percent in response to the shock. This corresponds to a

[15]Taken from https://www.bls.gov/opub/mlr/2015/article/labor-force-projections-to-2024-1.htm, last accessed on February 15, 2017.

2.56 percent decrease from the BAU growth rate, which means that the impulse response for output should start off at -2.56 percent in time period zero. Eventually, the growth rate following the shock converges back to the BAU growth rate of about 2.5 percent. This is fairly close to the 2015 growth rate of U.S. real GDP, calculated at 2.596 percent by the World Bank,[16] which indicates that my estimated VAR model and subsequent derivation have produced rather accurate forecasts.

Figure 5.2: Derived growth rates for output with and without an energy shock

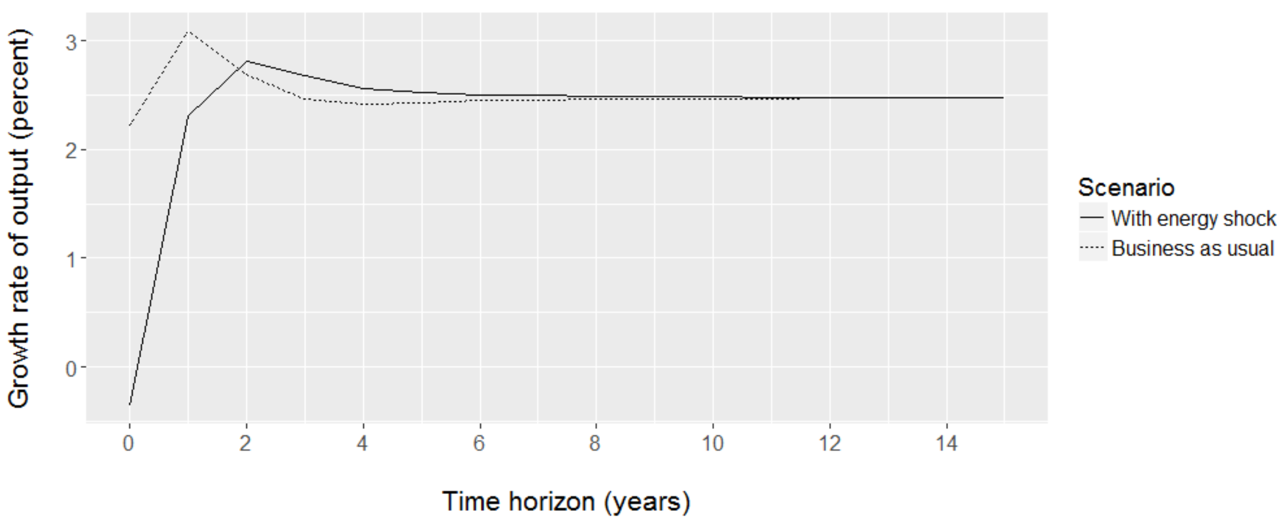

Next, I derive the impulse response for output directly from the impulse responses for the individual factor inputs using Equation 3.5. This is shown in Figure 5.3. The derived curve corresponds very well to the signed difference between the curve with the energy shock and the BAU curve in Figure 5.2. In particular, the impulse response starts off at -2.56 percent in year zero and tends to zero over the next 15 years. This gives a visual answer to my hypothesis about the long-run behavior of output following a negative shock to fossil fuel energy consumption. Although the growth rate decreases by as much as 2.56 percent

[16]Taken from http://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=US, last accessed on February 17, 2017.

in response to the shock, it quickly recovers to trend. In fact, as can be seen from the 95 percent upper confidence band, the output gap is statistically significant for only the first time period following the shock.

Figure 5.3: Derived impulse response for the growth rate of output (dashed lines give the bootstrapped 95 percent confidence intervals)

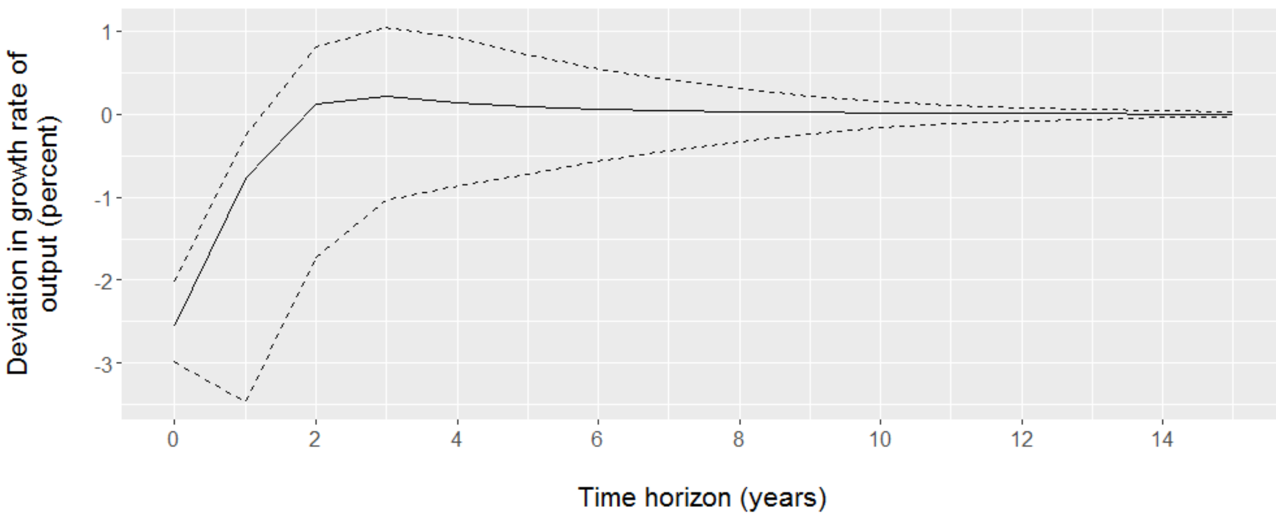

Now I perform a formal test of my main hypothesis. Taking 10 years after the shock to represent the long-run and letting $IRF_Y(t)$ denote the height of the impulse response function for output t years after the shock, the null and alternative hypotheses can be formulated as follows:

$$H_0 : IRF_Y(10) = 0 \qquad \text{i.e. growth returns to trend in the long-run}$$

$$H_A : IRF_Y(10) \neq 0 \qquad \text{i.e. growth does not return to trend in the long-run}$$

Assuming the null to be true, the test statistic $\frac{I\hat{R}F_Y(10)}{SE\left[I\hat{R}F_Y(10)\right]}$ approximately follows a standard normal distribution.[17] To compute the observed test statistic for this data, I refer to Table 5.2,

[17] The test statistic actually follows an asymptotic t-distribution with infinite degrees of freedom, which converges to the standard normal distribution.

which shows the estimated heights of the impulse response functions and the bootstrapped standard errors for selected years.[18] The observed test statistic is $\frac{0.0154}{0.221} = 0.0696$, which, when compared to the standard normal distribution, gives a two-tailed p-value of 0.945. There is insufficient evidence to conclude that the growth rate of output does not return to trend in the long-run.

Table 5.2: Estimated impulse response function for the growth rate of output (standard errors are bootstrapped standard errors)

Year	IRF_Y	$SE\left[IRF_Y\right]$
0	-2.563	2.006
1	-0.777	2.829
2	0.124	1.842
3	0.215	1.573
4	0.142	1.298
5	0.0858	1.018
6	0.0564	0.783
7	0.0401	0.577
8	0.0292	0.423
9	0.0213	0.307
10	**0.0154**	**0.221**
11	0.0112	0.159
12	0.00806	0.114
13	0.00582	0.0819
14	0.00421	0.0588
15	0.00304	0.0422

Bolded values are those that are important for my hypothesis test.

[18]Note that the standard errors for the impulse response for output cannot be obtained simply from linear combinations of those for the individual factor inputs, because the VAR variables are not statistically independent. Instead, Table 5.2 reports bootstrapped standard errors. Also note that the bootstrapped conditional distribution of the estimated heights given a time point is very asymmetric for time periods prior to 4 years, as shown in Figure 5.3, and so parametric hypothesis testing is only possible for the later years when the estimated heights become approximately normally distributed.

5.3 Critique of Model

Although the results successfully allowed me to test my main and two sub-hypotheses, the reduced form results in Table 5.1 still raise concerns about whether I have correctly specified my model. The low multiple R-squared values for ∇K_t and ∇E_t are driven in part by the lack of significance for the lagged versions of the same variables. I explore this problem by fitting univariate $AR(1)$ models for ∇K_t and ∇E_t separately.

The estimated $AR(1)$ model for ∇K_t has an insignificant slope term, which suggests that the growth rate of capital in the current year does not depend on the growth rate in the previous year. Because this would otherwise defeat the purpose of modeling capital using time series techniques, I checked the $AR(1)$ model for just the log of capital and found a significant slope term. Somehow, first-differencing capital had removed the lag one autocorrelation, such that the growth rate of capital might as well be white noise. However, I could not have modeled capital in log-levels because capital and fossil fuel energy were not co-integrated. The estimated $AR(1)$ model for ∇E_t has a significant slope term, which indicates that the growth rate of energy in the current year does depend on the growth rate in the previous year. Although this autocorrelation gets masked by the addition of other variables in the VAR, this is not a cause for concern as long as this inter-temporal dependency is being accounted for somehow. The potential multicollinearity in the estimating equation for ∇E_t is also not a problem because the focus of my analysis is not on factors that are important in forecasting energy consumption.

The low R-squared values can be improved by including more lags in the VAR model. To see how including more lags changes my results, I fit the same VAR model but this time with three lags, as recommended by half of the Information Criteria tests in Table 4.2. I find that the R-squared values rise marginally, from 0.0750 to 0.115 for ∇K_t and from 0.124 to 0.205 for ∇E_t. Although some of the lag two and three variables are significant, all but one of the lag one variables in Table 5.1 which were crucial for model-checking and hypothesis-testing lost significance. Additionally, the impulse responses no longer show statistically significant dynamic effects on TFP after the same 10 percent negative shock to fossil fuel energy consumption. Apparently, including three lags sacrifices too many degrees of freedom relative to the number of data points used. Because there is no strong theoretical justification for the inclusion of more than one lag, I conclude that the original model with only one lag is more appropriate for the data.

The exclusion of labor from the model has the potential to bias the impulse responses for the first few time periods after the shock. Because labor can be thought of as a complement to energy in the production process, a negative shock to energy consumption could conceivably cause a temporary contraction of the labor force. For example, higher carbon taxes would make it more expensive for firms to transport freight over long distances, leading to a decrease in the demand for truck drivers. The derived impulse response for output in Figure 5.3 should therefore see a larger initial decline to reflect the sudden retrenchment of workers. On the other hand, if labor and capital behave as economic substitutes, then the decline in labor could induce firms to hire more capital (Grandmont, Pintus and Vilder, 1998), canceling out the decrease in labor in the derived impulse response for output. Regardless, these short-run effects on labor should not affect the long-run growth rate of output, and so the answer to my hypothesis about the long-run growth rate returning back to trend does not change.

As a more general critique of VAR as an econometric tool, Stock and Watson (2001) point out that the impulse response functions tend to be very sensitive to the identifying assumptions used in the structural component, and hence the ordering of the variables. I perform additional sensitivity analysis by changing the ordering of TFP and energy (capital is still ordered first). Now, TFP is ordered second and energy is ordered third, and so sudden changes in fossil fuel energy consumption exert no contemporaneous effects on TFP. The impulse responses arising from this different set of identifying assumptions is shown in Figure 5.4. With the exception that TFP no longer falls instantaneously in year zero, these new impulse responses are very similar to those under the original set of identifying assumptions. The decline in TFP reaches a minimum of -0.033 (compared to the -0.044 from earlier) and remains statistically significant until about 5.5 years after the shock, while capital sees no statistically significant deviations from trend. The derived impulse response for output should therefore be similar to the one under the original set of identifying assumptions, particularly for the later time periods, and so the answer to my main hypothesis does not change. This shows that my model is robust to different identifying assumptions and that the concern put forth in Stock and Watson (2001) does not detract much from my results.

Figure 5.4: Impulse responses for a 10 percent negative shock to the growth rate of fossil fuel energy consumption, under the alternative set of identifying assumptions (dashed lines give the bootstrapped 95 percent confidence intervals)

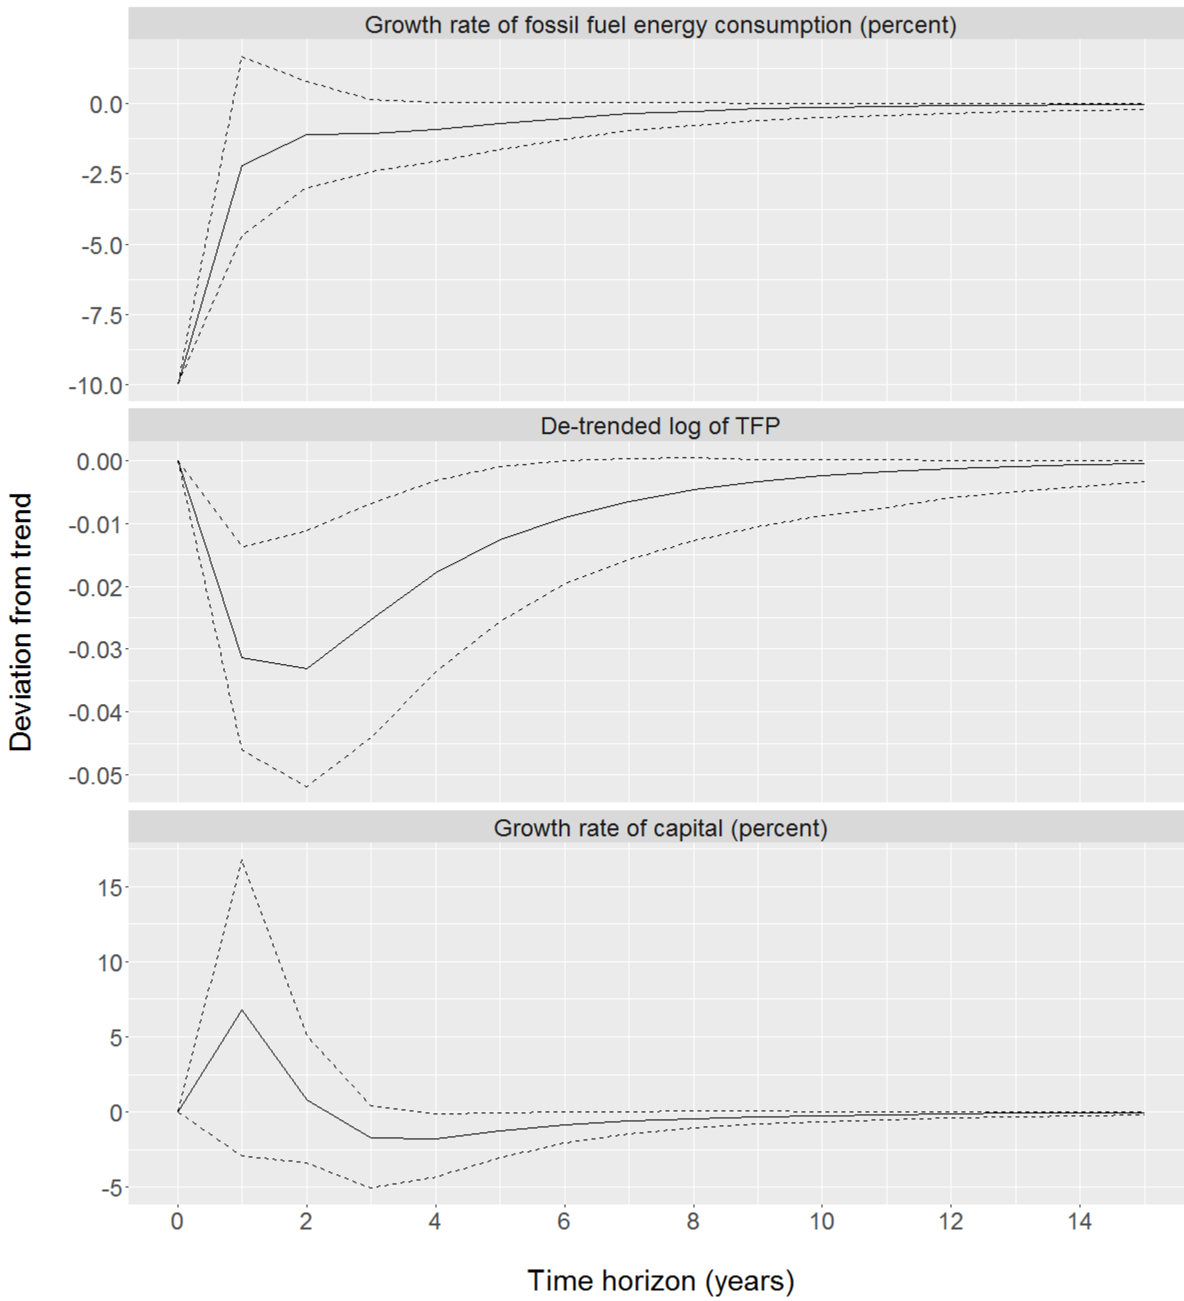

6. Conclusion

6.1 Implications for Climate Policy

To summarize my results, I found that a 10 percent negative shock to the growth rate of fossil fuel energy consumption initially causes the growth rate of output to fall by 2.56 percent but recover to trend quickly within as little as 1.5 years after the shock, such that the long-run growth rate remains unaffected. Moreover, the temporary output gap is driven more by productivity declines in the form of lower capital utilization and not by changes in the size of the capital stock, which saw no statistically significant changes due to the shock. These findings suggest three implications for climate policy.

First, climate policy that seeks to reduce emissions by cutting fossil fuel energy consumption results in output loss that is only temporary. This output gap is also very brief, lasting for no more than 1.5 years, and so there should be no worry that climate policy may stunt economic growth in the long-run. Second, this temporary output gap is not too large. The growth rate falls to a trough of -0.355 percent immediately after the 10 percent negative shock to energy, constituting a very mild recession. Of course, this output gap will be larger if I had included the short-run effects of the shock on labor, and will scale proportionately with the size of the initial shock to energy. Nonetheless, my analysis provides a useful approximation of the size of the output gap relative to that of the energy shock. Also, because statistical significance in the impulse responses is independent of the size of the initial shock, growth will still return to trend within 1.5 years regardless of how large the initial shock is. Lastly, the slump induced by climate policy should be cushioned by improving productivity to reduce the amount of idle capital, and not by increasing the amount of capital held by firms. Encouraging research in more energy-efficient production techniques, for example, would allow firms to continue running their physical capital at close to full capacity in the event of a policy-induced energy shock.

6.2 Further Work

Although this paper succeeds in understanding the long-run behavior of output, policy makers may also be interested in quantifying the short-run effects of an energy shock for cost-benefit analysis. The SVAR model developed in this paper is not yet precise enough for such a task. Most importantly, the size of the labor force should be included as an endogenous variable, because business cycles should be accounted for in any short-run analysis. Doing so, however, presents a highly intractable parameter identification problem involving four endogenous variables, and so future work should look into the theoretical literature for ways to impose the additional restrictions necessary to identify the SVAR system. Further research into the nature of capital and energy as substitutes or complements is also needed to determine *a priori* how the size of the capital stock should change following a negative energy shock, so that this may be used as a benchmark in model-checking. Whether capital and energy behave as substitutes or complements also clearly influences the size of the initial output gap.

Throughout this research, I restricted my analysis to just fossil fuel energy consumption to ignore possible substitution by firms among various energy sources. This is an oversimplification of reality, especially with the rise of non-polluting, renewable energy sources. Future work can study this switching behavior by incorporating consumption of non-polluting energy sources as yet another endogenous variable. After going through the parameter identification, one may find that a mandated cut to fossil fuel energy consumption causes firms to switch sufficiently to non-polluting sources such that the initial fall in output is cushioned to some extent. Hence, developing such a model is in the best interests of policy makers who want to show that even the short-run impacts of climate policy are not too great to bear.

Finally, future research should design similar SVAR models for data from developing countries, to which the concerns about stifled economic development apply most. Because the composition of the capital stock differs between developing and developed countries, the causality channels linking the endogenous variables may be different for the former. The impulse responses for a representative developing country may therefore differ wildly from those presented in this paper, possibly to the extent that idle capital can no longer be used to explain productivity declines following an energy shock. Of more value are comparative studies that contrast the effects of climate policy in developed and developing countries. Finding that a sudden reduction to energy consumption severely impairs long-run growth or induces a particularly long and hard recession in developing countries would mean that developing countries are less likely to pursue climate policy.

A. Bibliography

Ang, James B. "CO_2 emissions, energy consumption, and output in France." *Energy Policy* 35 (2007): 4772–4778.

Edelstein, Paul, and Lutz Kilian. "How sensitive are consumer expenditures to retail energy prices?" *Journal of Monetary Economics* 56 (2009): 766–779.

Ghali, Khalifa H., and M.I.T. El-Sakka. "Energy use and output growth in Canada: a multivariate cointegration analysis." *Energy Economics* 26 (2004): 225–238.

Golosov, Mikhail, John Hassler, Per Krusell and Aleh Tsyvinski. "Optimal taxes on fossil fuel in general equilibrium." (working paper, National Bureau of Economic Research, Cambridge, MA, 2011).

Gradus, Raymond, and Sjak Smulders. "The trade-off between environmental care and long-term growth – Pollution in three prototype growth models." *Journal of Economics* 58 (1993): 22–51.

Grandmont, Jean-Michael, Patrick Pintus, and Robin de Vilder. "Capital-labor substitution and competitive nonlinear endogenous business cycles." *Journal of Economic Theory* 80 (1998): 14–59.

Halicioglu, Ferda. "An econometric study of CO2 emission, energy consumption, income and foreign trade in Turkey." *Energy Policy* 37 (2009): 1156–1164.

Hassler, John, Per Krusell, and Conny Olovsson. "Energy-saving Technical Change" (working paper, National Bureau of Economic Research, Cambridge, MA, 2012).

Ivanov, Ventzislav, and Lutz Kilian. "A Practitioner's Guide to Lag Order Selection For VAR Impulse Response Analysis." *Studies in Nonlinear Dynamics & Econometrics* 9 (2005): 1–30.

Kim, Hyeongwoo. "Generalized Impulse Response Analysis: General or Extreme?" (working paper, Auburn University Department of Economics, 2012).

Koetse, Mark J., Henri L.F. de Groot, and Raymond J.G.M. Florax. "Capital-energy substitution and shifts in factor demand: A meta-analysis." *Energy Economics* 30 (2008): 2236–2251.

Ozturk, Ilhan. "A literature survey on energy-growth nexus." *Energy Policy* 38 (2010): 340–349.

Narayan, Paresh K., and Seema Narayan. "Carbon dioxide emissions and economic growth: Panel data evidence from developing countries." *Energy Policy* 38 (2010): 661–666.

Pao, Hsiao-Tien, and Chung-Ming Tsai. “CO_2 emissions, energy consumption and economic growth in BRIC countries.” *Energy Policy* 38 (2010): 7850–7860.

Phillips, P. C. B., and Ouliaris, S. “Asymptotic Properties of Residual Based Tests for Cointegration.” *Econometrica* 58 (1990): 165–193.

Prywes, Menahem. “A nested CES approach to capital-energy substitution.” *Energy Economics* 8 (1986): 22–28.

Shafik, Nemat. “Economic Development and Environmental Quality: An Econometric Analysis.” *Oxford Economic Papers, New Series* 46 (1994): 757–773.

Soytas, Ugur, Ramazan Sari, and Bradley T. Ewing. “Energy consumption, income, and carbon emissions in the United States.” *Ecological Economics* 62 (2007): 482–489.

Soytas, Ugur, and Ramazan Sari. “Energy consumption, economic growth, and carbon emissions: Challenges faced by an EU candidate member.” *Ecological Economics* 68 (2009): 1667–1675.

Stock, James H., and Mark W. Watson. “Vector Autoregressions.” *The Journal of Economic Perspectives* 15 (2001): 101–115.

Werf, Edwin van der. “Production functions for climate policy modeling: An empirical analysis.” *Energy Economics* 30 (2008): 2964–2979.

Winston, Gordon C. “The Theory of Capital Utilization and Idleness.” *Journal of Economic Literature* 12 (1974): 1301–1320.

Zhang, Xing-Ping, and Xiao-Mei Cheng. “Energy consumption, carbon emissions, and economic growth in China.” *Ecological Economics* 68 (2009): 2706–2712.