

Polgreen, Linnea; Silos, Pedro

Working Paper

Crude substitution: The cyclical dynamics of oil prices and the skill premium

Working Paper, No. 2006-14a

Provided in Cooperation with:

Federal Reserve Bank of Atlanta

Suggested Citation: Polgreen, Linnea; Silos, Pedro (2008) : Crude substitution: The cyclical dynamics of oil prices and the skill premium, Working Paper, No. 2006-14a, Federal Reserve Bank of Atlanta, Atlanta, GA

This Version is available at:

<https://hdl.handle.net/10419/70728>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Crude Substitution: The Cyclical Dynamics
of Oil Prices and the Skill Premium

Linnea Polgreen and Pedro Silos

Working Paper 2006-14a
August 2008

Crude Substitution: The Cyclical Dynamics of Oil Prices and the Skill Premium

Linnea Polgreen and Pedro Silos

Working Paper 2006-14a
August 2008

Abstract: Higher oil-price shocks benefit unskilled workers relative to skilled workers: At the business-cycle frequency, energy prices and the skill premia display a strong, negative correlation. We assess the robustness of this negative correlation using several methods and data sources, including sector-level data. We find that the negative correlation is robust to different de-trending procedures, and the wages of unskilled workers in energy-intensive industries have a larger positive correlation with oil prices. We also estimate the parameters of an aggregate technology, which uses, among other inputs, energy and heterogeneous skills. We find that both capital-skill and capital-energy complementarity are responsible for this correlation pattern. As energy prices rise, the use of capital decreases and the demand for unskilled labor relative to skilled labor increases, resulting in lower skill premia.

JEL classification: E24, E32, J24

Key words: skill heterogeneity, energy prices, business cycles, capital-skill complementarity

This paper is a comprehensive revision of a paper previously circulated under the same title. The authors thank Steven Durlauf, Nir Jaimovich, Karsten Jeske, Jim Nason, B. Ravikumar, Víctor Ríos-Rull, Ellis Tallman, Robert Tamura, and seminar participants at the Midwest Macroeconomic Meetings, Econometric Society Summer Meetings, Colgate and Stony Brook universities, from whom we have received many useful comments. They especially thank Martin Eichenbaum (the associate editor) and one anonymous referee for providing detailed comments that have greatly improved the paper. Finally, we also thank Mark Dumas and Steve Rosenthal from the U.S. Bureau of Labor Statistics for help with the KLEMS dataset. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Linnea Polgreen, College of Pharmacy, University of Iowa, Iowa City, IA 52242, 319-335-3797, 319-335-1956 (fax), linnea-polgreen@uiowa.edu, or Pedro Silos, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309, 404-498-8630, 404-498-8956 (fax), pedro.silos@atl.frb.org.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's Web site at www.frbatlanta.org. Click "Publications" and then "Working Papers." Use the WebScriber Service (at www.frbatlanta.org) to receive e-mail notifications about new papers.

1 Introduction

Both oil prices and the skill premia¹ have been increasing for the past forty years. Despite this long-term, overall increasing trend, when examined closely, over shorter time periods, oil prices and the skill premia fluctuate in opposite directions, displaying a very strong, negative correlation. This negative correlation is specifically found at the business-cycle frequency, and it is robust to different detrending methods.

To examine and quantify the mechanism that leads to the negative correlation between oil prices and the skill premia, we estimate an aggregate-production function with an explicit role for energy, and conclude that capital-skill complementarity² and capital-energy complementarity are responsible for this correlation pattern. Due to capital-energy complementarity, a rise in energy prices decreases the amount of capital used. Capital-skill complementarity increases the demand of unskilled labor (relative to skilled labor), decreasing the skill premia.

The paper has two parts. The first provides a detailed analysis of the data. In the second part we use these data to estimate a structural model. Using annual energy-price and skill-premium data for the last four decades, we assess the robustness of their unconditional correlation to different de-trending procedures. Specifically, we filter the data in three different ways and consistently find that this unconditional correlation is negative and statistically significant. Using unconditional correlations has two disadvantages. First, if the original data are not stationary, they need to be transformed to induce stationarity. Different procedures to do so might lead to different results. Second, focusing only on unconditional correlations can lead to a loss of useful information about the joint dynamics of the two variables. Thus to further analyze the comovement between the skill premia and oil prices, we compute a set of statistics proposed by Den Haan (2000) to measure the comovement between output and consumer prices. This involves computing

¹The skill premium is the ratio of the wages of the college-educated to those without a college degree.

²Capital-skill complementarity is the idea that capital is more complementary with skilled rather than unskilled labor.

the correlation coefficient of the individual forecast errors implied by a vector autoregression at each of several different forecast horizons. As the choice of which variables to include in the VAR could potentially alter the results, we experiment with different VAR specifications. We begin by fitting a VAR to skill premium and oil prices only. We find that correlations are significantly negative for almost all forecast horizons, except the one-period-ahead forecast. We assess the sensitivity of these results to two additional VAR specifications. The first includes a measure of exogenous energy-price movements, which we model as a dummy variable taken from Hoover and Perez (1992). The second modification includes both output and a short-term nominal interest rate as endogenous variables.

Finally, we analyze data for different sectors. Assuming frictions to inter-sectoral labor mobility are small, one would expect to find a stronger relationship between the skill premia and energy prices in energy-intensive industries, given that all other technology parameters are constant. We group industries by their degree of energy intensity using data from the Multi-Factor Productivity program of the Bureau of Labor Statistics. We compare skill premia and wages for different skill types in energy-intensive industries to those in non-energy-intensive industries. A graphic representation of the time series data of filtered skill premia and oil prices illustrates that the skill premia in energy-intensive sectors react more strongly to changes in energy prices.

The second part of the paper tests the validity of the hypotheses of capital-skill and capital-energy complementarities, by specifying a five input aggregate production function (including energy) and estimating its parameters. We use aggregate data on capital equipment, non-residential structures, labor inputs for different skill types, and energy use and prices to estimate the production function in its original non-linear form. This exercise can be viewed as extending Krusell *et al.*'s (2000) analysis to a framework in which energy use and prices are explicitly introduced. Based on our parameter estimates, we find strong evidence of capital-skill complementarity, although the estimated elasticities do not differ much quantitatively from those found with similar data sets but without energy

in the production function. We also find evidence of capital-energy complementarity. Moreover, the correlation between the de-trended fitted skill premia and oil prices is of same magnitude as that observed in the data.

(Literature review/context in which this paper stands– to be added)

2 Energy Prices and the Skill Premium

The skill premium is a weighted ratio of skilled wages to unskilled wages³. We define skill by education level: a skilled worker has a college degree, and an unskilled worker does not. Data are obtained from the Current Population Survey (CPS), 1963–2004.

Data on energy prices and usage come from the US Government Energy Information Administration. We use annual data from 1949 to 2004 for prices and quantities of oil, coal and natural gas, which represent almost 85% of overall energy consumption in the US. The price of energy used throughout the analysis is a Laspeyres index of the prices of those three main energy sources. The final energy price index was the result of dividing the constructed energy price index by the Gross Domestic Product deflator.

Using an index of oil, coal and natural gas allows us to compare our results to previous studies, such as Kim and Loungani (1992). Because oil is a large percentage of total energy consumption in the US economy, the deviation from trend of the constructed price index has a very large correlation (about 0.98) with the deviation of oil prices. If oil prices were used instead of the measure used here, all results presented would still hold.

Deviations of the energy price index relative to its HP-trend are shown in Figure ??.

Nominal energy prices were very stable until about 1974, decreasing sharply relative to overall inflation. The first oil shock occurs in 1974, when prices rose 78%. The second major oil price increase occurs five years later, during 1979 and 1980, when prices increased by 27% and 35%, respectively. Large oil price increases did not occur for the next 20 years. However, in 1999, oil prices went up by 21% and by 41% in 2000. After the first two oil crises, there were two large price drops, occurring in 1985 (-52%) and in 1998 (-30%).

³Details are provided in the appendix.

Thus, overall, energy prices have displayed a large amount of volatility over the last three decades.

2.1 Oil Prices and the Skill Premium: Unconditional Correlations

Figures 1 - 3 show the de-trended skill premium and energy prices, using three types of de-trending methods: deviations from an exponential trend, a band-pass filter⁴ that removes fluctuations occurring in periods smaller than three or larger than 35; and a (log)HP-filtered series with a smoothing parameter equal to 100. We report correlation coefficients for the three de-trending procedures (in parentheses, we also report standard errors) in Table 1.⁵ Correlations are negative and in some cases surprisingly strong. For instance, the correlation coefficient between oil prices and the skill premium after removing an exponential trend is -0.72, with a standard error of only 0.07. It is still strong using the band-pass filter and somewhat weaker using the HP filter. With the latter de-trending procedure a two-standard-deviation interval does not include zero, but it is close.⁶

Table 1: Correlations: Skill Premium and Oil Prices

Filter	Entire Sample	Second Subsample
Exp. Detrend	-0.713 (0.066)	-0.690 (0.091)
BP Detrend	-0.476 (0.119)	-0.505 (0.140)
HP Detrend	-0.312 (0.154)	-0.343 (0.173)

On the second column of Table 1 we report the same correlation but assuming that the data begin in 1979. We eliminate the first oil shock and the large drop in the skill premium in the mid-seventies. As the reader can see, correlations move very little.

⁴We use the band-pass filter proposed by Christiano and Fitzgerald (2004).

⁵Standard errors were computed from an exactly identified GMM procedure. Estimates of the first and second moments were estimated using moment conditions, using a weighting matrix proportional to the covariance matrix of the residuals. By the Delta Method we compute standard errors for the correlation coefficients, for which case the analytical gradient has a simple expression.

⁶The MATLAB function `corrcoef.m` provides probability values for testing the hypothesis of no correlation. The band-pass and exponential detrending are significant at the 1% level. The HP-filtered series is significant at the 5% level.

2.2 Evidence from Vector Autoregressions

In the introduction we outlined two drawbacks of focusing only on unconditional correlations. First, the analyst has many ways of inducing stationarity in the two series and results might not be robust across different methodologies. Second, unconditional correlations could lead to a loss of valuable information about the dynamic relationship between two variables. With this in mind, in this section we compute a set of statistics proposed by Den Haan (2000) to measure the degree of comovement between consumption prices and output. More specifically he suggests examining the correlation coefficients at different forecast horizons using the moments implied by the estimated coefficients of a VAR. Given a vector of N endogenous random variables Y_t and M exogenous random variables X_t we fit the following model,

$$Y_t = \alpha + \gamma t + BX_t + \sum_{l=1}^L A_l Y_{t-l} + \epsilon_t \quad (1)$$

where A_l are $N \times N$ matrices, α and γ are vectors of dimension N and B is an $N \times M$ matrix. The vector of innovations ϵ_t is also of dimension N , it has a mean of zero and a covariance matrix Ω . These innovations are serially uncorrelated but they can be correlated with each other.

Our interest in the relationship between oil prices and the skill premium implies that Y_t needs to include, at least, these two variables. We restrict L , the number of lags, to be one as our sample is rather short, and we experiment with three specifications. The first VAR includes only oil prices and the skill premium in the endogenous vector and no exogenous variables X_t . In the second specification we include, as an exogenous variable, the dates that Hoover and Perez (1992) identified as purely exogenous oil prices shocks that were triggered, for example, by the Suez episode (1956-57) or the Iranian revolution (1978-79). These dates (omitting the month) are: 1947, 1953, 1957, 1969, 1970, 1974, 1978, 1979, 1981, and they enter the VAR as a variable that takes the value of 1 if there was an oil price shock that year and 0 otherwise. The inclusion of these exogenous shocks

could imply a different correlation structure once their effect has been accounted for. Finally, a third specification includes four variables in the endogenous vector: oil prices, skill premium, real output, and the nominal rate of return of a 3-month Treasury bill. In this third specification we omit the Hoover-Perez dummy variable.⁷ One advantage of the methodology suggested by Den Haan (2000) is that it avoids the need to impose stationarity in the data. The different choices for doing so can lead to different results. For all three specifications, we estimate the VAR with non-stationary data as we only take logarithms of the endogenous variables without any additional filtering.

We sample from the posterior distribution of the VAR parameters using non-informative priors, which amounts to sampling from the VAR’s likelihood function. For each draw of the posterior ⁸, we compute the correlation of the forecast errors of oil prices and the skill premium at different forecast horizons. We plot the median of this correlation coefficient along with the 10th and 90th percentiles in the three panels of Figure 4. Each panel shows results for one of the three different specifications: oil prices and skill premium only; the Hoover-Perez dummy as an additional exogenous variable; and interest rates and real output included in the endogenous vector.

2.3 Evidence from Sectoral Data

We further investigate these correlations by examining sectoral data. Whatever relationship exists between oil prices and the skill premium, it is likely to be stronger in industries where energy is a relatively more important input. We obtained data from the Multi-Factor Productivity program of the Bureau of Labor Statistics (BLS). These data provide annual measures of outputs and inputs for 20 manufacturing industries from 1949 to 2001. The measure of output is the total value of production measured in cur-

⁷The results were robust to including the Hoover-Perez dummy and to first-difference output prior to estimation.

⁸We used well-know methods for sampling from the posterior distribution of the parameters. In particular, we used a Gibbs sampler for multivariate regressions as described in Zellner (1974, pp. 224-227) We sampled a total of 5,000 draws discarding the first 1,000 before computing the statistics of interest.

rent dollars. For inputs, these data include expenditures (in current dollars) for capital, labor, energy, materials, and purchased business services. All these measures are flows; for example, expenditures in capital equal an estimated service flow from an industry's capital stock. Expenditures for labor are measured as paid hours of work in a particular industry.⁹ For additional details on how these measures are constructed the reader is referred to Gullickson and Harper (1987).

The first step in our analysis is to separate the 20 industries into two groups: energy-intensive industries and non-energy-intensive industries. To measure energy intensity we focus on the ratio of energy expenditures to total value of production (E/VP). We compute this ratio as a time-series average of the annual expenditures in energy to value of production. We label an industry as energy intensive when its E/VP is larger than the manufacturing-industry average. The first column of Table 2 displays the value of this ratio for the 20 manufacturing industries, the durable goods industries, the non-durable goods industries, and total manufacturing. Classifying industries as energy-intensive as those with a E/VP ratio larger than the average of the entire manufacturing sector, gives us five energy-intensive industries: textiles, chemicals, paper, primary metals, stone and glass. Defining energy intensity using alternative ratios does not completely alter the set of industries that belongs to the energy-intensive group. We compute two ratios in the last two columns of Table 2: energy-to-capital expenditures and energy-to-capital-plus-labor expenditures. According to the second ratio, food and rubber and plastics should be included in the energy-intensive group. According to the third ratio, food, petroleum refining, and rubber and plastics would be considered energy-intensive. The five industries initially identified as energy intensive can be classified as such irrespective of the ratio used. As a result, we define these industries as the energy-intensive group.

⁹The energy expenditures include energy used for heat and power only. If fuel is a good transformed in the manufacturing process, it will not be part of the energy expenditures category. For example, petroleum refining entails transforming an intermediate good (crude oil) into a final good (refined oil). Purchases of crude oil for this industry would be included in materials expenditures and not in energy expenditures.

Table 2: Energy Intensity Across Manufacturing Industries

Industry	E/VP	E/K	E/(K+L)
Tobacco	0.006	0.028	0.017
Apparel	0.008	0.116	0.019
Printing/Publishing	0.008	0.064	0.015
Transp. Equipment	0.009	0.179	0.020
Misc. Manufacturing	0.011	0.090	0.022
Furniture	0.011	0.149	0.025
Electronics	0.011	0.073	0.019
Measuring instruments	0.011	0.113	0.022
Ind./Comm. Machinery	0.011	0.091	0.021
Leather	0.012	0.121	0.022
Food	0.015	0.162	0.052
Fabricated Metal	0.015	0.139	0.033
Petrol. Refining	0.021	0.120	0.084
Lumber/Wood	0.022	0.131	0.039
Rubber/Plastics	0.022	0.249	0.053
Textiles	0.035	0.333	0.071
Chemicals	0.052	0.224	0.099
Paper	0.055	0.332	0.111
Primary Metals	0.060	0.634	0.135
Stone/Glass	0.070	0.577	0.133
Durables	0.025	0.168	0.040
Non Durables	0.027	0.161	0.057
Total Manufacturing	0.027	0.159	0.046

We use information in the CPS to identify the industry in which individuals work, and we group them as belonging to either the energy-intensive sector or the non-energy-intensive sector. We then compute the skill premium analogously to the aggregate case. The sample sizes we work with are not large (around 3800 individuals per year for the energy-intensive industries and 7600 for the non-energy-intensive). Consequently, to remove some noise associated with our samples we compare the two sectors using band-pass filtered estimates of prices and skill premia.¹⁰ This removes some higher-frequency movements that exponential de-trending and HP-filtering retain. In addition, we report covariances and correlation coefficients. When standard deviations for the two series dif-

¹⁰The band pass filter smooths further the higher frequencies relative to the filter used in section 2.1. Here we remove frequencies higher than five periods.

fer, and they do differ for the skill premia in our two groups of industries, covariances can be a better measure of comovement. And it is stronger co-movement that a larger energy share predicts. Table reports covariances and correlation coefficients along with standard errors.¹¹

Table 3: Oil Prices and Skill Premia Across Sectors

Sector	$Covar(SP_t, P_t)$	$Corr(SP_t, P_t)$
Energy Intensive	-0.009 (0.002)	-0.614 (0.080)
Non Energy Intensive	-0.008 (0.002)	-0.675 (0.073)

The table shows that the point estimate for the covariance is larger in absolute value for the energy intensive industries (-0.009 vs. -0.008), implying a greater degree of comovement. As a result of the short time series, there is substantial overlap of the distributions of the two point estimates: the covariance for the energy-intensive sector is within one standard error of that of the non-energy-intensive sector, so the two do not differ statistically. The correlation coefficient is a bit larger in absolute value for the sector in which energy is less important, but this is a result of a larger volatility in the non-energy-intensive sector.

Table 4: Oil Prices and Unskilled Wages Across Sectors

Sector	$Covar(W_{u,t}, P_t)$	$Corr(W_{u,t}, P_t)$
Energy Intensive	0.0014 (0.0017)	0.1057 (0.1235)
Non Energy Intensive	-0.0013 (0.0013)	-0.1118 (0.1087)

Table 5: Oil Prices and Skilled Wages Across Sectors

Sector	$Covar(W_{s,t}, P_t)$	$Corr(W_{s,t}, P_t)$
Energy Intensive	-0.0074 (0.0023)	-0.4083 (0.0943)
Non Energy Intensive	-0.0089 (0.0025)	-0.4726 (0.0958)

¹¹These standard errors are computed in the same way as those in section 2.1.

Figure 5 plots band-pass-filtered oil prices and skill premia for the two groups of industries. The energy-intensive sector reacts more strongly, in general, to oil price movements: the drop in prices at the beginning of the seventies, the rise (and minor fall) during the seventies, the fall during the mid-nineties, and the recent drop in the face of rising energy prices during the first years of the current decade. The only exception is the rise in energy prices around 1990, in which the skill premium in the non-energy-intensive sector dropped more than in the energy-intensive sector.

Visual inspection of the series points to a more pronounced response of the skill premium in the energy-intensive sector to changes in oil prices. However, by analyzing the relationship between skill premia and oil prices across the two groups of industries, we have found no statistically significant differences between measures of comovement. Correlation coefficients estimates are larger (in absolute value) in the non-energy-intensive sectors, while the estimated covariances point to a stronger relationship between energy prices and the skill premium in the energy-intensive sector. As a result of the short time series, the distributions for these estimates overlap substantially.

Given economic theory, however, the limited difference between these two sectors is not surprising: this sectoral analysis is predicated on three fragile assumptions. First, assuming that energy prices have a larger impact on the skill premium in industries with a larger share of energy assumes frictions in inter-sectoral labor mobility. From a theoretical perspective, insofar as labor is perfectly mobile across sectors within a period, wages would equalize (for a given skill level) in all sectors. The skill premium, in turn, would react to a change in energy prices in every industry in exactly the same manner. Second, it also assumes that other technology parameters are constant across industries. In particular, differences in elasticities of substitution across sectors would have an obvious impact on the effects of energy prices on the skill premium. Unfortunately, we have no ex-ante information about these differences. Finally, when we focus on just energy shares (measured by E/VP), differences across industries are not large. The share in textiles, an energy-intensive industry, is 3.5% which is close to the average for the entire manufacturing

sector (2.7%).

3 Estimation of an Aggregate Production Function

An explanation for the previous correlation patterns demands a structural estimation of an aggregate production function. Our hypothesis of capital-skill complementarity and capital-energy complementarity, which would lead to the observed correlation, needs to be tested. We do so by specifying an aggregate technology for the U.S. economy and estimating its parameters.

The theoretical model to be estimated is derived from a profit-maximizing firm's first-order conditions for choosing from among five factors of production: skilled labor (s_t), unskilled labor (u_t), structures (k_{st}), energy (e_t), and equipment (k_{et}). The production-function form combines a CES aggregation of unskilled labor, an aggregation of equipment and energy (the capital-energy composite), and an aggregation of skilled labor and the capital-energy composite. This aggregate combines with structures through a Cobb-Douglas function:

$$Y_t = G(e_t, k_{st}, k_{et}, u_t, s_t) = k_{st}^\alpha [\mu u_t^\sigma + (1-\mu)(\lambda(\xi k_{et}^\nu + (1-\xi)e_t^\nu)^{\frac{\rho}{\nu}} + (1-\lambda)s_t^\rho)^{\sigma/\rho}]^{(1-\alpha)/\sigma}, \quad (2)$$

where μ , λ , and ξ are parameters that govern income shares, and σ , ρ , and ν are parameters that drive the elasticities of substitution between equipment and unskilled workers, equipment and skilled workers, and energy and equipment respectively. The firm purchases capital equipment units at a (per unit) price q_t , energy units at a price p_t , and units of structure at a (normalized) price of unity. Energy and equipment prices follow stochastic processes known by the firm owner. Moreover, factor markets are assumed to be perfectly competitive. The firm can rent equipment units at a rental rate equal to r_t . Finally, purchased units of capital equipment and structures depreciate at rates δ_e and δ_s respectively.

We define the elasticity of substitution between the energy-equipment composite and

unskilled labor, the energy-equipment and skilled labor, and energy and equipment to be $\frac{1}{1-\sigma}$, $\frac{1}{1-\rho}$, and $\frac{1}{1-\nu}$ respectively.¹² In addition, the skilled and unskilled labor inputs, s_t and u_t are functions of hours (h_s and h_u) and efficiency indices (ψ_s and ψ_u): $s_t = \psi_{st}h_{st}$ and $u_t = \psi_{ut}h_{ut}$.

Defining by G_{i_t} as the marginal product of input i at time t , the first order conditions for profit-maximizing firm imply the following equations:

$$p_t = G_{e_t} \quad (3)$$

$$w_{s,t} = G_{h_{s,t}} \quad (4)$$

$$w_{u,t} = G_{h_{u,t}} \quad (5)$$

$$r_t = G_{h_{u,t}} \quad (6)$$

$$\frac{q_{t-1}}{q_t} = \frac{1}{(1 - \delta_e)} \{(1 - \delta_s) - G_{k_{st}} - q_{t-1}G_{k_{et}}\} + \epsilon_t \quad (7)$$

The first four equations equate rental rates to marginal products for four different inputs: energy, skilled labor, unskilled labor, and equipment capital. The last equation is a no-arbitrage condition that sets the expected return on equipment equal to the expected return on structures, where ϵ_t is an equipment-price-forecast-error which is normally distributed with a mean of zero and a variance equal to σ_ϵ^2 .

The estimation is done in two steps. In the first step, we only estimate the parameter driving the elasticity of substitution between energy and capital equipment, ν . In the second step, we estimate all the remaining parameters of the model. The reason to separate the estimation into two different parts is that estimating ν can be done by OLS using a very simple structural relationship. The second step in the estimation is much more involved.

¹²In defining these as the elasticities of substitution we are assuming that no other factors change except the pair of factors that we are considering. When the number of inputs in production is only two this is not an issue. However, in production technologies with more than two inputs there are several ways one can define the elasticity of substitution between any pair while accounting for changes in all other inputs. Two widely used measures are the Allen and the Morishima elasticities. Please see Polgreen and Silos (2008) for a discussion in the context of a similar model and for additional references.

Dividing equation (6) by equation (3), we obtain,

$$\frac{r_t}{p_t} = \frac{G_{ket}}{G_{et}} = \frac{\xi}{1 - \xi} \frac{k_{et}^{\nu-1}}{e_t^{\nu-1}} \quad (8)$$

A straightforward manipulation gives,

$$\frac{r_t k_{et}/Y_t}{p_t e_t/Y_t} = \frac{G_{ket}}{G_{et}} = \frac{\xi}{1 - \xi} \frac{k_{et}^\nu}{e_t^\nu} \quad (9)$$

The left hand side is the ratio of capital's share to output and the ratio of energy expenditures to output. Denote this left-hand side variable as $rkey_t$. The right hand side is a constant times the ratio of capital to energy raised to ν . Denoting the ratio of capital to energy as rke_t , and taking logs and first differences yields,

$$\Delta \log(rkey_t) = \nu [\Delta \log(rke_t)] \quad (10)$$

The parameter ν can be estimated consistently by regressing $rkey_t$ on rke_t and we describe the construction of these two series in an Appendix. Figure 6 displays them from 1950 to 2004. The dashed-dotted line is $rkey_t$ and the solid line rke_t . Ordinary least-squares estimation gives a value for ν of -0.962 with a standard error of 0.461. The case of $\lim_{\nu \rightarrow 0}$ corresponds to a Cobb-Douglas aggregate between energy and equipment, so $\nu = -0.962$ implies substantially more complementarity; the elasticity of substitution is only about 0.5. We use this estimate of ν to fix its value in the second part of the estimation, which we now describe.

Manipulating optimality conditions (5) and (4) gives us the following two equations:

$$\frac{w_{st} h_{st} + w_{ut} h_{ut}}{Y_t} = (1 - \alpha) \{ \mu (h_{ut} \psi_{ut})^\sigma + (1 - \mu) [\lambda (\xi k_{et}^\nu + (1 - \xi) e_t^\nu)^\frac{\rho}{\nu} + (1 - \lambda) (h_{st} \psi_{st})^\rho]^\sigma \}^{-1} \{ \mu (h_{ut} \psi_{ut})^\sigma + (1 - \mu) [\lambda (\xi k_{et}^\nu + (1 - \xi) e_t^\nu)^\frac{\rho}{\nu} + (1 - \lambda) (h_{st} \psi_{st})^\rho]^\frac{\sigma}{\rho} \}^{-1} (1 - \lambda) (h_{st} \psi_{st})^\rho \}, \quad (11)$$

and,

$$\frac{w_{st}}{w_{ut}} = \frac{1 - \mu}{\mu} \sigma (1 - \lambda) [\lambda (\xi k_{et}^\nu + (1 - \xi) e_t^\nu)^\frac{\rho}{\nu} + (1 - \lambda) (h_{st} \psi_{st})^\rho]^\frac{\sigma}{\rho} \frac{(h_{st} \psi_{st})^\rho}{(h_{ut} \psi_{ut})^\sigma} \frac{h_{ut}}{h_{st}}, \quad (12)$$

Equation (11) equates the share of labor in output to a non-linear function of exogenous variables, latent variables, and parameters. The left-hand side variable of equation (12) is the skill premium.

The no-arbitrage condition (7) equates the growth rate of the relative price of capital equipment to a non-linear function of parameters, exogenous variables, and latent variables. We can stack conditions (7),(11), and (12) to yield the following equation,

$$W_t = f(\theta; X_t, \psi_{st}, \psi_{ut}, \epsilon_t) \quad (13)$$

Here W_t is the vector of left-hand side variables: the share of capital in output, the skill premium, and the growth rate of equipment prices. We group the set of exogenous variables $h_{st}, h_{ut}, e_t, k_{et}, k_{st}$ in the vector X_t . The vector θ includes all the unknown parameters in the model, $(\xi, \nu, \sigma, \rho, \mu, \lambda, \alpha, \delta_e, \delta_s, \sigma_\epsilon^2)'$.

The sources of estimation errors are given by the price-forecast-error ϵ_t and the latent variables ψ_{st} and ψ_{ut} which follow the stochastic process,

$$\phi_t = \phi_0 + \nu_t, \quad (14)$$

where $\phi_t = [\log(\psi_{st}), \log(\psi_{ut})]'$ and $\nu_t \sim N(0, \Sigma)$. The covariance matrix Σ is diagonal and we restrict the two diagonal elements to be equal to σ_ψ^2 .

Equations (13) and (14) are the measurement equations and transition equations of a non-linear state-space model.¹³ There are several methods one can use to estimate its parameters and latent variables, but we choose a Bayesian procedure employed by Polgreen and Silos (2008).¹⁴ Bayesian inference in our environment involves specifying a prior distribution $p(\gamma)$ for the parameters of interest $\gamma = (\theta, \Sigma, \phi_0)$; constructing a posterior distribution $p(\gamma | \{W_t\}_{t=1}^T, \{X_t\}_{t=1}^T)$ as the product of the prior and the likelihood

¹³For technical reasons we need to include additive measurement errors in the first two equations of W_t . The variances of these errors turn out to be small.

¹⁴A complete description of the estimation methodology is outside the scope of this paper. The interested reader is referred to Polgreen and Silos (2005) for a detailed description of the procedure. For alternative methodologies see Ohanian *et al.* (2000).

function $L(\{W_t\}_{t=1}^T|\gamma, \{X_t\}_{t=1}^T)$. Any statistics of interest can be obtained by sampling from the posterior distribution.¹⁵

3.1 Priors

For most of the parameters, we use the same prior distributions as those used by Polgreen and Silos (2008). Following Krusell *et al.* (2000), to aid in the estimation we fix three of those parameters: the two depreciation rates δ_e and δ_s , and ν , the parameter driving the elasticity of substitution between capital and energy, which we estimated above. The depreciation rates for equipment and structures were fixed at 0.1250 and 0.005. ν is fixed at -0.962. Energy introduces an additional parameter ξ which we endow a prior normal distribution with mean 0.5 and standard deviation 0.1, truncated to the $[0, 1]$ region. Table 6 summarizes our priors.

Table 6: Prior Distributions

Parameter	Prior
ξ	$N(0.5, 0.1) \chi_{[0,1]}(\xi)$
σ	$N(0.575, 0.25) \chi_{[-\infty, 1]}(\sigma)$
ρ	$N(-0.76, 0.25) \chi_{[-\infty, 1]}(\rho)$
μ	$N(0.5, 0.2) \chi_{[0,1]}(\mu)$
λ	$N(0.5, 0.2) \chi_{[0,1]}(\lambda)$
σ_ϵ^2	$Gamma(0.3, 0.01)$
α	$N(0.11, 0.005) \chi_{[0,1]}(\xi)$
σ_ψ^2	$Gamma(0.4, 0.01)$

The prior mean for ρ is halfway between 0.08, estimated by Berndt and White (1978), and -1.6 estimated by Dennis and Smith (1978). These studies cover the manufacturing sector from 1950-1973. The prior mean for σ is the same as the estimate from Clark and Freeman (1977), and a number also reported in Hammermesh's (1993) survey of labor demand. The share parameters μ , λ , and ξ have prior distribution centered at the mid-point of their admissible regions, with relatively large standard deviations. The prior

¹⁵The results we present below are based on 300,000 draws from the posterior distribution.

on α , the share of structures, has a rather informative prior, given its minor role in the analysis. Its prior mean is centered at Krusell *et al.*'s estimate, which in turn is close to the value calibrated by Greenwood, Hercowitz and Krusell (1997) and equal to 0.13. Priors on the variances are relatively diffuse.

3.2 Estimation Results

Table 7 reports posterior means and standard deviations for σ and ρ . We also include the estimate for ν , obtained above by OLS. The first two parameters drive the elasticities of substitution of equipment with unskilled and skilled labor respectively; ν drives the elasticity of substitution between energy and capital equipment. The posterior moments for ρ and σ are close to those obtained by Polgreen and Silos (2008); see their Table 1, third line. At their means the estimates for ρ and σ imply values for the elasticities of substitution between equipment with unskilled and skilled labor equal to 4.4 and 0.65, respectively. These estimates imply a large degree of capital-skill complementarity; while the previously found estimate of ν implies equipment-energy complementarity.

Table 7: Posterior Moments

Parameter	Posterior Mean (or OLS estimate)	Posterior Standard Deviation (or OLS s.e.)
σ	0.774	0.045
ρ	-0.525	0.066
ν	-0.962	0.461

With the draws from the posterior distribution of the parameters we can readily obtain a “distribution” for the fitted skill premium resulting from this model. We construct it the following way. We first set all shocks in the model to zero at all points in time. We then use each draw and the value of the exogenous variables (capital, hours, etc. . .) to construct a fitted skill premium for our sample period using the right hand side of equation (12). We de-trend each of these fitted values using the three procedures used in Section 2.1. For each draw of the posterior and for each of the de-trending methods we can compute

a (posterior) correlation coefficient between the skill premium and oil prices. Once we have this entire distribution we can obtain from it any statistic of interest. We report posterior means and standard deviations of this distribution of correlation coefficients on the first column of Table 8. The table shows how the fitted skill premium is negatively

Table 8: Fitted Skill Premium vs. Oil Prices

De-trending Proc.	Posterior Distribution	GMM s.e.'s
Exp. De-trending	-0.736 (0.019)	-0.814 (0.048)
BP Filter	-0.544 (0.048)	-0.648 (0.081)
HP Filter	-0.189 (0.009)	-0.349 (0.114)

correlated with oil prices irrespective of the methodology one uses to de-trend. These (mean) negative correlations are sufficiently far away from zero and of a similar magnitude as those found with actual data. A possible exception is the HP-detrended skill premium, which has a weaker correlation with oil prices than that observed with actual data. This is a consequence of the model's inability to capture the really high-frequency component of the skill premium. However, we find these results quite remarkable given the absence of technology shocks in the model.

Finally, to further compare our results to those found in Section 2.1 we also compute standard errors using that same GMM procedure. We compute once the fitted skill premium with the posterior means of the parameters and the exogenous variables. We emphasize again that we "turn off" all shocks in the model for all time periods. We compute the correlation of the skill premium and oil prices and compute the GMM-standard-errors (again, for each the three procedures) and we report them on the second column of Table 8. These magnitudes suggest an even stronger relationship between the skill premium and oil prices than that observed in the data. Notice that all estimated correlations are closer to -1 than with actual data, except perhaps when we HP-detrend, in which case the magnitude is about the same.

4 Conclusion

The relative wage that a skilled worker earns relative to that earned by an unskilled worker, the skill premium, is negatively correlated with oil prices at the business cycle frequency. Here, we have clearly established the robustness of this fact. We employed three de-trending methods (an HP filter, a band-pass filter, and deviations from an exponential trend), three VAR model specifications (one with oil prices and the skill premium as endogenous variables, one including exogenous oil shocks, and another using short-term-interest rates and output), and data disaggregated by the intensity of oil used in each industry group. Even with all these permutations, in all cases except one (a one-period ahead VAR), we find evidence for the negative correlation between energy prices and the skill premium.

In addition, we have estimated an aggregate production function in which energy use and prices are explicitly introduced. In the estimation we obtain two key results. First, capital is more substitutable with unskilled labor than with skilled labor. However, this is not controversial: a wide body of research has found some degree of capital-skill complementarity in the US economy (e.g., Griliches (1969), Krusell *et al.* (2000)). Also, capital-skill complementarity has been used to explain the low frequency movements of the skill premium (e.g., Krusell *et al.* (2000)). Second, there is a high degree of complementarity between capital and energy. These two facts are a plausible explanation for the observed correlation between oil prices and the skill premium: when oil prices rise, firms substitute unskilled workers for capital, and the skill premium falls.

A Data

Skilled and Unskilled Wages: The skill premium is calculated using a method from Polgreen and Silos (2005). We obtain data from the CPS (March out-going rotation) and include anyone who is at least 16 but not over 70 years old. We include only those who have wage and salary income. (This excludes the self employed.) Many observations have missing hours: the CPS asks what one's income was last year, but how many hours one worked last week. Thus, interviewees who were on vacation or on any other type of leave during the previous week would have income from last year, but no hours for last week. In order to retain as many observations as possible, we impute missing hours. Hours are estimated using age, age², years of education and dummy variables representing female, black, and white. We then eliminate any observation that is missing any necessary variable or has unreasonable hourly wages¹⁶.

The hours variable is then multiplied by the number of weeks worked last year to obtain annual hours. The annual hours, l , are weighted by the CPS weights, μ , and an ability index, $w_{g,96}$, representing the average wage in 1996 of similar individuals, g . Annual hours, l , are summed over all observations in each skill level, j , to obtain the labor input series¹⁷, N .

$$N_{j,t} = \sum_{i \in G_{j,t}} l_{i,t} w_{g,96} \mu_{i,t},$$

where i represents each observation, and t represents the year. The wage is calculated by multiplying the wage, $w_{i,t}$, by the annual hours variable and the CPS weights, summing over all observations for each skill level, and dividing by the labor input series.

$$W_{j,t} = \frac{\sum_{i \in G_{j,t}} w_{i,t} l_{i,t} \mu_{i,t}}{N_{j,t}}.$$

¹⁶Following Card and DiNardo 2002), we consider unreasonable wages to be less than \$1 or greater than \$100 in 1979 dollars.

¹⁷The sample is divided into 264 groups based on age, race, gender and education level, and we calculate the average wage of each group in 1996 to create an ability index. To make the index unitless and to avoid problems with inflation, the index is then divided by the average wage in 1996 for each skill level. This is an appropriate ability index: if one's wage represents one's marginal product, those with higher wages represent a larger amount of labor input per hour. See Denison (1979).

The numerator is the total wage bill: the average wage in the group times the average labor input in the group, weighted by the CPS weights. This is divided by the labor input, N , to get the wage series for both the skilled and the unskilled. The skill premium is then calculated by dividing the wage series for the skilled by the wage series for the unskilled.

We transform nominal wages for a given group by the consumption expenditures deflator in the National Income and Product Accounts (NIPA), Haver Analytics code:?????.

Capital Stock Data: Data on capital stocks comes from the annual estimates provided by Haver Analytics. For private structures we take the series ????, and for private equipment and software we take the series ????. We assume that the price of a unit of structures is unity and the price of a unit of equipment is q_t . We therefore deflate the stock of structures by the consumption implicit price deflator (series ????) and the stock of equipment and software by the implicit price deflator of investment in equipment and software (series ????). The relative price q_t is the ratio of the two deflators, the investment deflator divided by the consumption deflator.

Energy Use: For our structural estimation energy use enters in units of BTU. We compute the total units consumed in the U.S. economy of oil, natural gas, and coal. To estimate ν , the first step in the structural estimation, we construct the variable ke_t as the ratio of the deflated capital equipment stock (see previous paragraph) to the units of energy use.

Capital and Labor Share in Output: To compute capital's share of output we take the sum of (nominal) Rental Income, Corporate Profits, Net Interest, and Capital Depreciation and we divide it by nominal GDP minus Proprietors' income. This assumes that the share of capital in Proprietors' income is the same as the share in the overall economy. Except GDP all variables come from Table 1.12, (NIPA).

To compute the share of labor in output, we take the ratio of Wage Compensation to private employees, and we divide it by nominal GDP minus Proprietors' Income. Again, these variables also come from Table 1.12 in NIPA.

Constructing $rkey_t$: The left hand side variable in the regression equation (10) is capital's share in output divided by the ratio of total (nominal) expenditures in energy to nominal GDP. Nominal expenditures in energy come from the EIA website.

References

- [1] REFERENCES TO BE ADDED.

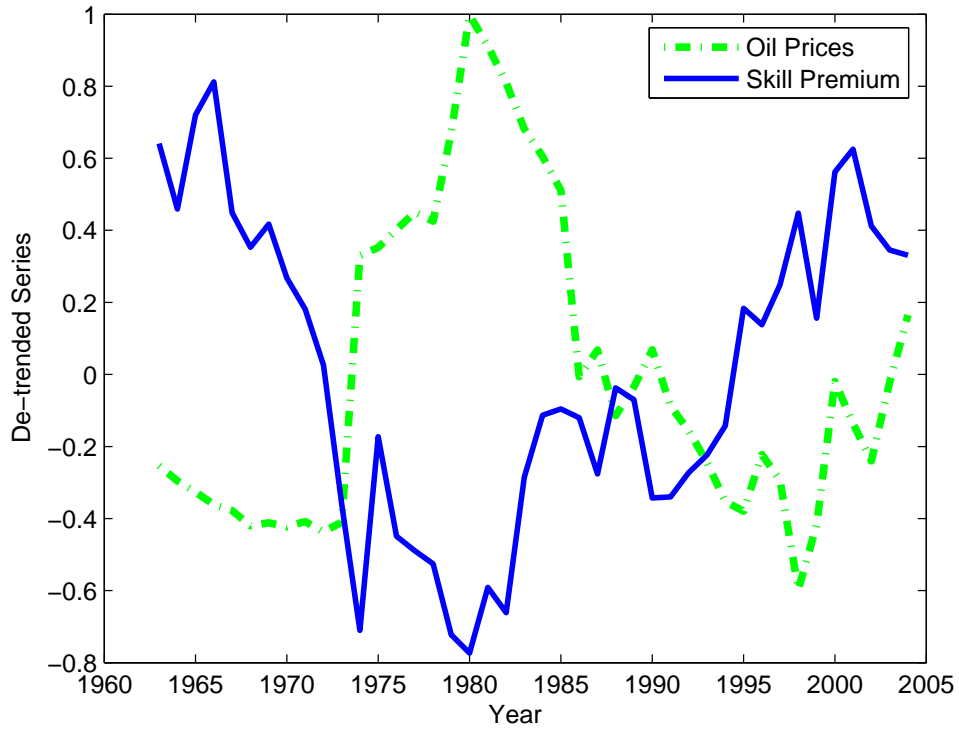


Figure 1: Exp.Detrend

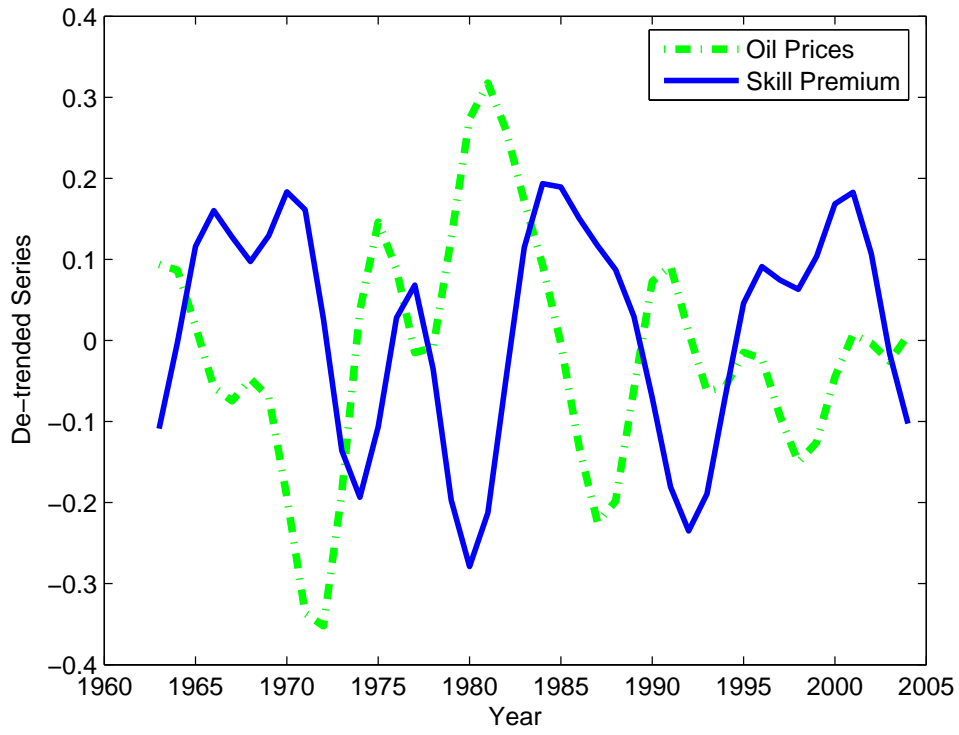


Figure 2: Exp.Detrend

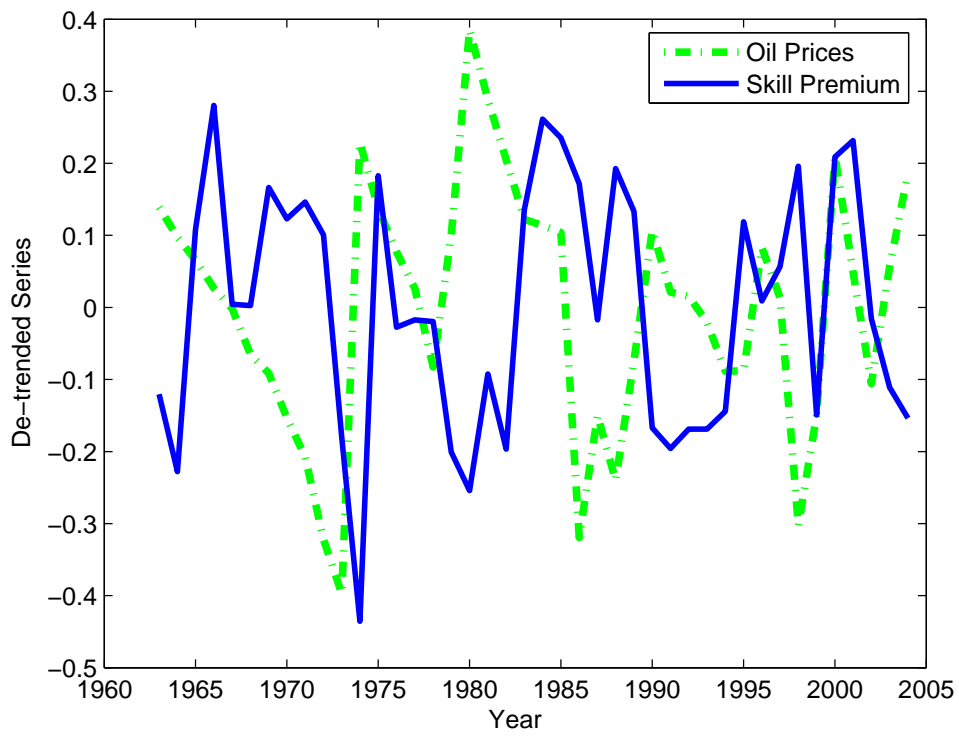


Figure 3: Exp.Detrend

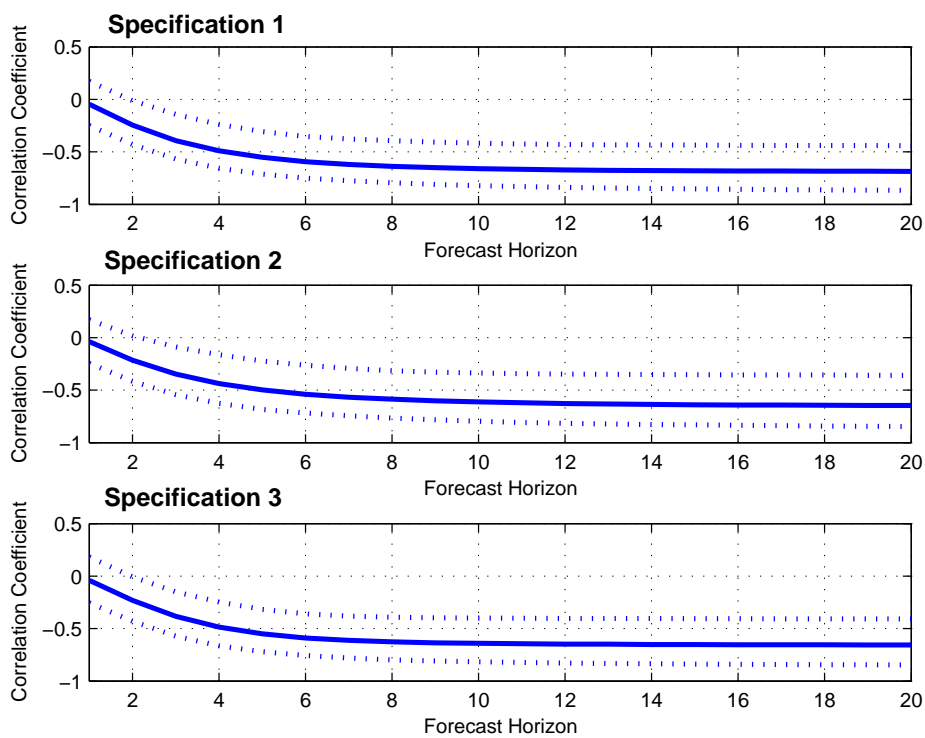


Figure 4: Correlation coefficient of forecast errors of the skill premium and oil prices at different forecast horizons for different VAR specifications. The solid line plots the median of the distribution of the correlation coefficient at a particular horizon and the dotted lines plot the 10th and 90th percentiles of that distribution. Specification 1 includes oil prices and the skill premium as endogenous variables. Specification 2 adds the Hoover-Perez dummy as an exogenous variable. Specification 3 includes a nominal short term interest rate and real output as endogenous variables.

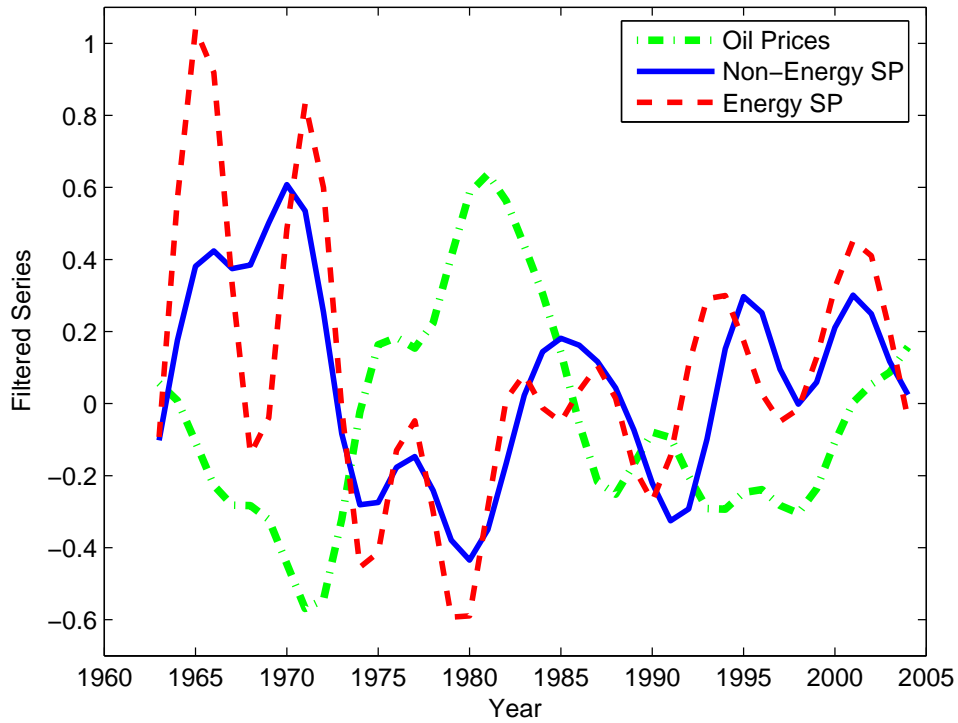


Figure 5: Band-pass filtered series of oil prices and skill premia by industry group. The dashed-dotted line (left axis) is the filtered oil price series; the dashed line (right axis) is the skill premium in the energy intensive sector; and the solid line the skill premium in the non energy-intensive sector.

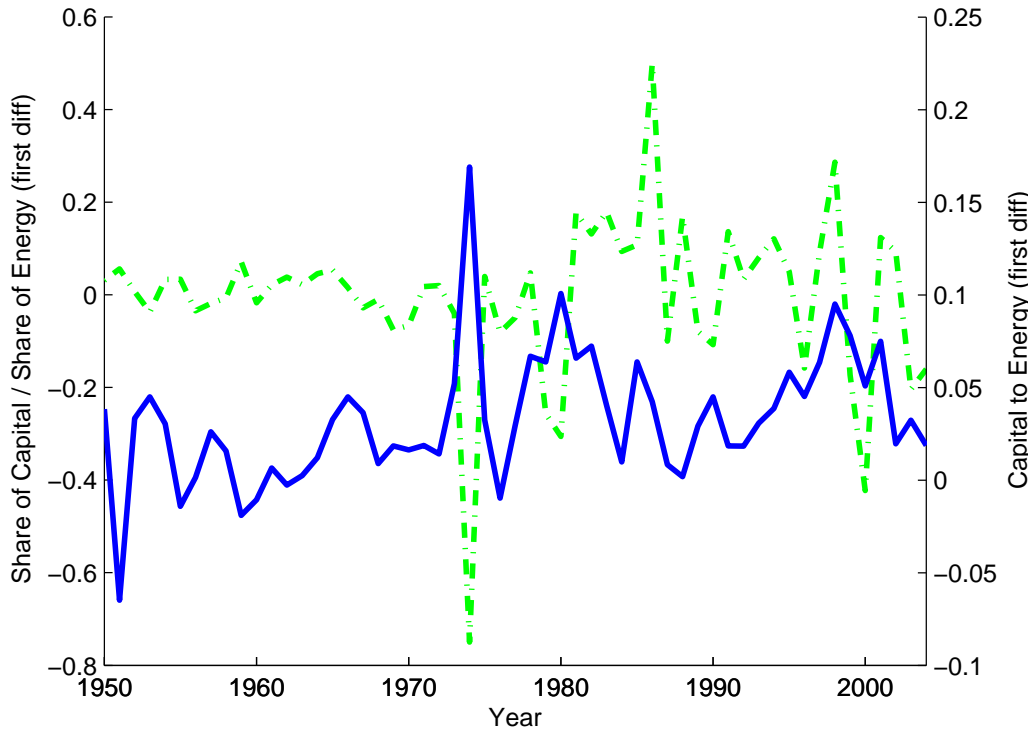


Figure 6: Two series to estimate ν . The dashed-dotted line (left-hand axis) is the ratio of the share of capital in output to the share of energy in output. The solid line (right-hand axis) is the ratio of capital to energy.