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

Linnea Polgreen*†
University of Iowa

Pedro Silos‡
Federal Reserve Bank of Atlanta

February 23, 2009

Abstract

At the business cycle frequency, energy prices and the skill premium display a strong, negative correlation. This fact is robust to different de-trending procedures. Identifying exogenous shocks to oil prices using the Hoover-Perez (1992) dates, shows that the skill premium falls in response to such a shock. The estimation of the parameters of an aggregate technology that uses, among other inputs, energy and heterogeneous skills, demonstrates that capital-skill and capital-energy complementarity are responsible for this correlation. As energy prices rise, the use of capital decreases and the demand for unskilled labor – relative to skilled labor – increases, lowering the skill premium.

Keywords: skill-heterogeneity, energy prices, business cycles, capital-skill complementarity

JEL Classification: E24, E32, J24

*This is a comprehensive revision of a paper previously circulated under the same title. Thanks to Steven Durlauf, Nir Jaimovich, Karsten Jeske, Federico Mandelman, Jim Nason, B. Ravikumar, Víctor Ríos-Rull, Ellis Tallman, Robert Tamura, and seminar participants at the Midwest Macroeconomic meetings, Econometric Society summer and winter meetings, and Colgate and Stony Brook Universities, from whom we have received many useful comments. We 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 Bureau of Labor Statistics for help with the KLEMS dataset used in previous versions of the paper. A Data Appendix describing the sources and the construction of all series used in the paper, as well as replication programs and data sets, is available at the journal website.

†College of Pharmacy, University of Iowa, Iowa City, IA 52242, 319-335-3797 (phone), 319-335-1956 (fax), linnea-polgreen@uiowa.edu.

‡Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree St. NE, Atlanta, GA 30309, 404-498-8630 (phone), 404-498-8956 (fax), Pedro.Silos@atl.frb.org.

1 Introduction

Over the past four decades and at the business cycle frequency, oil prices and the skill premium fluctuate in opposite directions displaying a very strong, negative correlation. This pattern is robust to different de-trending methods.

To examine and quantify the mechanism that leads to the negative correlation between oil prices and the skill premium, this paper estimates an aggregate production function with an explicit role for energy and conclude that capital-skill complementarity – the idea that capital is more complementary with skilled rather than unskilled labor– 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 premium.

This paper has two parts. The first provides a detailed analysis of the data. In the second part, these data are used to estimate a structural model. Using annual energy-price and skill premium data for the past four decades, this paper assesses the robustness of the unconditional correlation between oil prices and the skill premium to different de-trending procedures. Specifically, filtering the data in three different ways consistently shows that this unconditional correlation is negative and statistically significant. To overcome potential endogeneity problems, this analysis moves beyond unconditional correlations and estimate the response of the skill premium to an exogenous change in oil prices. These exogenous movements are isolated using a methodology proposed by Edelberg, Eichenbaum, and Fisher (1999) and Ramey and Shapiro (1998). These authors estimate the response of several macroeconomic aggregates to an exogenous change in government expenditures. They do so by identifying events – arguably independent of US economic conditions – that led to large military buildups. The dates of these events are in turn included in a VAR as an exogenous variable, making the response of the endogenous variables to the onset of one of these events easy to compute. Analogously, our analysis uses the Hoover and Perez (1992) dates for political events in the Middle East that disturbed oil production or expectations of oil production. Estimating the response of the skill premium, oil prices, and other variables of interest to the occurrence of a Hoover-Perez event, shows that oil prices rise and the skill premium falls. The fall in the skill premium is significant for a period of about three years, and it is robust to several VAR specifications.

The second part of the paper tests the validity of the hypotheses of capital-skill and capital-energy complementarities. It does so by specifying a five input aggregate production function (including energy) and estimating its parameters. Using aggregate data on capital equipment, nonresidential structures, labor inputs for different skill types, and energy use and prices, this paper estimates 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. Our parameter estimates imply a strong degree of capital skill complementarity, although the estimated elasticities do not differ significantly or quantitatively from those found with similar data sets but without energy in the production function. They also imply capital-energy complementarity. Moreover, the correlation between the de-trended fitted skill premium and oil prices is of same magnitude as that observed in the data.

Previous researchers have largely ignored energy prices in the study of the skill premium. To our knowledge, this paper is the first to empirically document this fact at the aggregate level and examine the relationship between cyclical movements in the skill premium and oil prices within a (partial) equilibrium model. Although work focusing on the behavior of the skill premium in equilibrium models does exist (e.g. Krusell *et al.* (2000) and Lindquist (2004)), energy use and prices and their implications for inequality are absent. Only one paper has specifically examined the effect of oil prices on relative wages: Keane and Prasad (1996) developed an empirical model using panel data and found that skilled, rather than unskilled workers, gain during oil price increases. Our line of work is different in a substantive way. First, this paper provides a structural interpretation of the data based on our estimates of the different elasticities of substitution. Second, it also provides a detailed analysis of the facts based on different data sources and methods.

2 Energy Prices and the Skill Premium: The Facts

The skill premium is a weighted ratio of skilled wages to unskilled wages[1]. Our definition of skill is by education level: a skilled worker has a college degree, and an unskilled worker does not. Data come from the Current Population Survey (CPS), 1963–2004.

Data on energy prices and usage come from the US Government Energy Information Administration. The time series of prices and quantities of oil, coal and natural gas, which represent almost 85% of overall energy consumption in the United States cover the period 1949 to 2004. 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.

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

[1] Details are provided in the data appendix available at the journal website.

0.98) with the deviation of oil prices. If oil prices were used instead of the measure here all results presented would still hold, and therefore the terms energy and oil prices are used interchangeably.

2.1 Oil Prices and the Skill Premium: Unconditional Correlations

The three panels of Figure 1 show the de-trended skill premium and energy prices[2], using three types of de-trending methods: deviations from an exponential trend, a band-pass filter [3] that removes fluctuations occurring in periods smaller than 3 or larger than 35; and a (log) HP-filtered series with a smoothing parameter equal to 100. Table 1 reports correlation coefficients for the three de-trending procedures (in parentheses, it also reports standard errors).[4] 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 -0.71, 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.[5]

The second column of Table 1 reports the same correlation but assumes that the data began in 1979, thus eliminating the first oil shock and the large drop in the skill premium that occurred in the mid-seventies. The changes in the correlation coefficients are small.

[2]Because energy prices are much more volatile than the skill premium, in all plots the skill premium is "magnified" by multiplying it by 10.

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

[4]Standard errors are computed from an exactly identified GMM procedure. Estimates of the first and second moments are estimated using moment conditions with a weighting matrix proportional to the covariance matrix of the residuals. By the Delta Method, the standard errors for the correlation coefficients are computed, in which case the gradient has a simple expression.

[5]The 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 The Response to an Exogenous Oil Price Shock

Unconditional correlations can mask an endogenous response of both oil prices and the skill premium to a change in US economic conditions. In this case, the argument of a re-allocation of factor inputs in response to a change in input prices as an explanation for the observed negative correlation between the skill premium and oil prices, would cease to be valid. Ideally, one would isolate the exogenous component of oil prices and would test whether the skill premium indeed fell in response to a rise in that component. Arguably, a large part of exogenous movements in oil prices are related to political instability in the Middle East, which is independent of US economic conditions. Using an indicator variable for the occurrence of a political event that disrupted oil production in the past is therefore a way to identify, if not all, at least the bulk of those exogenous oil price changes. Hoover and Perez (1992) constructed such indicator variable, and we label a Hoover-Perez episode the dates of those political events that caused large swings in energy prices.[6] This section investigates the response of the skill premium and other variables of interest to the onset of a Hoover-Perez episode by estimating a VAR in which those dates appear as an exogenous variable. This analysis is reminiscent of Ramey and Shapiro (1998) and Edelberg, Eichenbaum, and Fisher (1999), who identify exogenous increases in public expenditures by isolating events – which were independent of US economic conditions – in which large military buildups took place. Here, the Hoover-Perez dates play an analogous role to the Ramey-Shapiro episodes in those works. The fitted models are of the form:

$$logX_t = \alpha + AlogX_{t-1} + BHP_t + \epsilon_t. \tag{1}$$

[6] Following Bernanke, Gertler, and Watson (1997), the August 1990 invasion of Kuwait by Iraq is included as an additional Hoover-Perez episode.

The number of lags in the VAR is restricted to be one as our time series is rather short. In the previous equation, X_t is an $m \times 1$ vector of endogenous variables, α is an $m \times 1$ vector of constants, A is an $m \times m$ matrix, HP_t is a the date-t value of the Hoover- Perez variable (1 if t is a Hoover-Perez date, and zero otherwise), B is an $m \times 1$ vector of coefficients, and finally, ϵ_t is a zero-mean $i.i.d$ process with covariance matrix Ω. The response of $logX_{t+h}$ to a change in the value of HP_t is given by the coefficient on L^h in the polynomial $(I - AL)^{-1}B$.

An important modeling choice is which variables to include in the endogenous vector X_t. Our initial bivariate specification includes only the skill premium and oil prices. Including oil prices is important as this VAR approach would provide no basis for our analysis if it was found that oil prices failed to increase after a Hoover-Perez event. The two graphs in Figure 2 show the response of the two elements of $logX_t$ – (log) oil prices (top graph) and the (log) skill premium (bottom graph) – over a period of 15 years to a unit-change in HP_t. The solid, thicker line shows the median response and the two dotted lines show 66% confidence bands.[7] As expected, oil prices rise after a Hoover-Perez episode, with the effects peaking immediately and lasting for approximately 12 years. The skill premium falls after a Hoover-Perez event, and the median response remains negative for about 9 years (but it is only significant for 3). As with oil prices, the peak response happens immediately after the episode.

Given that large changes in energy prices have been associated with recessions in the

[7]The computation of these error bands uses the same bootstrapping procedure as the one described in Edelberg, Eichenbaum, and Fisher (1999). Specifically, given a vector $\{\hat{\epsilon_t}\}_{t=1}^{T}$ of fitted residuals from the VAR, one can sample with replacement from that vector to generate an artificial series $\{\tilde{\epsilon_t}\}_{t=1}^{T}$. Using the initial conditions and the estimated parameters of the fitted VAR, one can simulate an artificial series of the endogenous variable $log(X_t)$. Re-estimating the VAR using this new simulated series, one can compute the impulse responses in the same way as with the original data. Repeating this procedure 500 times, sorting the responses for each horizon by size, and taking the 17th, 50th, and the 83rd percentiles, yields the median response and the lower and upper bands.

United States, a VAR that also includes a nominal variable and a measure of real output was fitted. This VAR includes the Consumer Price Index (CPI) as our nominal variable and Real Gross Domestic Product (Real GDP) as a proxy for real output. The resulting vector of endogenous variables, $logX_t$, includes four time series: the log of the CPI, the log of Real GDP, the log of oil prices, and the log of the skill premium. The four panels of Figure 3 display the response of the four elements of $logX_t$ over 15 years to the occurrence of a Hoover-Perez event. The top two graphs show the response of oil prices and the skill premium and the bottom two the response of real output and the CPI. The inclusion of a measure of output and a nominal variable reduces the impact of the Hoover-Perez dates on oil prices, resulting in a more muted response. However, oil prices still rise after a Hoover-Perez date, and the peak effects are felt immediately. The response is significantly positive for 3 years. The response of the skill premium is roughly unchanged both qualitatively and quantitatively: the drop is significant for approximately 3 years, and the response peaks in the year of the event. The median response of output is negative for the first years but zero is within the error bands, illustrating the results of much empirical work that stresses the weakening relationship between oil prices and US output. Finally, there is a strong association between rises in oil prices and rises in inflation, and consequently, the response in the price index is positive – and very significant – for approximately 8 years.

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. This section specifies an aggregate technology for the US economy and estimates 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:

$$Y_t = G(e_t, k_{st}, k_{et}, u_t, s_t) = k_{st}^{\alpha}[\mu u_t^{\sigma} + (1-\mu)(\lambda \tilde{k}_t^{\rho} + (1-\lambda)s_t^{\rho})^{\sigma/\rho}]^{(1-\alpha)/\sigma}, \tag{2}$$

and,

$$\tilde{k}_t = (\xi k_{et}^{\nu} + (1-\xi)e_t^{\nu})^{\frac{1}{\nu}}, \tag{3}$$

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.

The elasticities of substitution between the energy-equipment composite and unskilled labor, the energy-equipment and skilled labor, and energy and equipment are given by $\frac{1}{1-\sigma}$,

$\frac{1}{1-\rho}$, and $\frac{1}{1-\nu}$ respectively.[8] 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 $u_t = \psi_{ut} h_{ut}$.

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

$$p_t = G_{e_t} \tag{4}$$

$$w_{s,t} = G_{h_{s,t}} \tag{5}$$

$$w_{u,t} = G_{h_{u,t}} \tag{6}$$

$$r_t = G_{h_{u,t}} \tag{7}$$

$$\frac{q_{t-1}}{q_t} = \frac{1}{(1-\delta_e)} \left\{ (1-\delta_s) - G_{k_{st}} - q_{t-1} G_{k_{et}} \right\} + \epsilon_t \tag{8}$$

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. The first step only estimates the parameter driving the elasticity of substitution between energy and capital equipment, ν. The second step estimates all the remaining parameters of the model. The reason to separate the estimation

[8] In defining these as the elasticities of substitution underlies the assumption that no other factors change except the pair of factors under consideration. 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.

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. Throughout it is assumed that variables chosen by the firm, and therefore endogenous, – $k_{et}, k_{st}, e_t, h_{ut}, h_{st}$ – are taken as exogenous by the econometrician. These variables are labelled observed independent variables.

Dividing equation (7) by equation (4), yields

$$\frac{r_t}{p_t} = \frac{G_{k_{et}}}{G_{e_t}} = \frac{\xi}{1-\xi}\frac{k_{et}^{\nu-1}}{e_t^{\nu-1}} \tag{9}$$

A straightforward manipulation gives

$$\frac{r_t k_{et}/Y_t}{p_t e_t/Y_t} = \frac{G_{k_{et}}}{G_{e_t}} = \frac{\xi}{1-\xi}\frac{k_{et}^{\nu}}{e_t^{\nu}} \tag{10}$$

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 yield

$$\Delta log(rkey_t) = \nu[\Delta log(rke_t)] \tag{11}$$

The parameter ν can be estimated consistently by regressing $rkey_t$ on rke_t, and the Appendix describes the construction of these two series. Figure 4 displays the series from 1950 to 2004. The dashed and dotted line is $\Delta log(rkey_t)$ and the solid line is $\Delta log(rke_t)$. Ordinary least-squares estimation gives a value for ν of -0.962 with a standard error of 0.461. The case of $\lim_{\nu\to 0}$ 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. The parameter ν is fixed at this value in the second part of the estimation, which is described below.

Let us denote by X_t the set of observed independent variables h_{st}, h_{ut}, e_t, k_{et}, and k_{st}, and by θ the vector of all unknown parameters parameters in the model, $(\xi, \nu, \sigma, \rho, \mu, \lambda, \alpha, \delta_e, \delta_s, \sigma_\epsilon^2)'$. Manipulating optimality conditions (5) and (6) gives us the following two equations:

$$\frac{w_{st}h_{st} + w_{ut}h_{ut}}{Y_t} = f_1(\theta; X_t, \psi_{st}, \psi_{ut}, \epsilon_t), \tag{12}$$

and,

$$\frac{w_{st}}{w_{ut}} = f_2(\theta; X_t, \psi_{st}, \psi_{ut}, \epsilon_t). \tag{13}$$

Equation (12) equates the share of labor in output to a non-linear function of observed independent variables, latent variables, and parameters. The left-hand side variable of equation (13) is the skill premium.

The no-arbitrage condition (8) equates the growth rate of the relative price of capital equipment to a non-linear function of parameters, observed independent variables, and latent variables. Stacking conditions (8),(12), and (13) yields the following equation,

$$W_t = f(\theta; X_t, \psi_{st}, \psi_{ut}, \epsilon_t) \tag{14}$$

Here W_t is the vector of left-hand side variables: the share of labor in output, the skill premium, and the growth rate of equipment prices.

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,

$$\phi_t = \phi_0 + \upsilon_t, \tag{15}$$

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

Equations (14) and (15) are the measurement equations and transition equations of a non-linear state-space model.[9] One can use several methods to estimate its parameters and latent variables, but we choose a Bayesian procedure employed by Polgreen and Silos (2008).[10] Bayesian inference in our environment involves specifying a prior distribution $p(\gamma)$ for the parameters of interest $\gamma = (\theta, \Sigma, \phi_0)$, and 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 function $L(\{W_t\}_{t=1}^T|\gamma, \{X_t\}_{t=1}^T)$. We can then obtain any statistics of interest by sampling from the posterior distribution.[11]

3.1 Priors

For most of the parameters, prior distributions are the same as those used by Polgreen and Silos (2008). Besides fixing ν at the value estimated above, the two depreciation rates δ_e and δ_s are fixed as well, following Krusell *et al.* (2000). 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, endowed with a prior normal distribution with mean 0.5 and standard deviation 0.1, truncated to the $[0, 1]$ region. Table 2 summarizes our priors.

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-

[9] The inclusion of additive i.i.d. measurement errors in the first two equations of W_t is done for technical reasons. The variances of these errors turn out to be small.

[10] A 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).

[11] The results presented below are based on 300,000 draws from the posterior distribution.

point of their admissible regions, with relatively large standard deviations. The prior on α, the share of structures, is rather informative, 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 3 reports posterior means and standard deviations for σ and ρ. It also includes 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.

With the draws from the posterior distribution of the parameters, one can readily obtain a "distribution" for the fitted skill premium resulting from this model. Its construction is as follows. First, all shocks in the model are set to zero at all points in time. One can 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 (13). These fitted values are de-trended using the three procedures in Section 2.1. For each draw of the posterior and for each of the de-trending methods, one can compute a (posterior) correlation coefficient between the skill premium and oil prices. Once this distribution is obtained, it is straightforward to compute any statistic of interest. Table 4, in its first column, reports posterior means and standard deviations of this distribution of correlation coefficients.

The table shows how the fitted skill premium is negatively 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. An exception is the HP-de-trended skill premium, which has a weaker correlation with oil prices than that observed with actual data. The weaker correlation is a consequence of the model's inability to capture the really high-frequency component of the skill premium.

To further compare our results to those found in Section 2.1, standard errors were computed using that same GMM procedure. Using the posterior means of the parameters and the exogenous variables, and "turning off" all shocks in the model for all time periods, the fitted skill premium was computed once. Table 4 reports on its second column the correlation of oil prices and the fitted skill premium along with its GMM-standard-error (again, for each the three procedures). 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 with HP-de-trending, in which case the magnitude is about the same.

In US data, oil prices are much more volatile than the skill premium. The ratio of the standard deviation of oil prices to the standard deviation of the skill premium in the United States is 9.852 (0.916), if one uses exponential de-trending; 10.510 (1.327) if one uses a band-pass filter; and 9.966 (1.205) if one uses an HP-filter. Table 5 is analogous to Table 4, but instead of displaying correlation coefficients, it displays the ratio of the standard deviation of oil prices to the standard deviation of the fitted skill premium. This ratio is roughly in line with the data for two of the de-trending procedures – band-pass and HP filters – with a value of approximately 8. If one uses exponential de-trending, the volatility of oil prices relative to the fitted skill premium is substantially lower – roughly half – than with actual skill premium.

Oil price shocks have been associated with recessions in the United States, particularly those of the 1970s and 1980s. Consequently, it is informative to compare the output and skill premium joint dynamics in the model with those found in US data. In particular, the focus is on the ratio of output and skill premium volatilities and the cross-correlations between output and the skill premium at one lead and one lag,[12] which Table 6 reports.

In the data, the volatilities between the skill premium and output are roughly the same. The point estimate of σ_{GDP}/σ_{SP} is about 1.16, but the standard error is 0.16, so a reasonable confidence band should include one. In terms of dynamic correlations, the skill premium leads the cycle, and the contemporaneous correlation is close to zero, 0.21, with a standard error of 0.15. Turning to the predictions from our model, Table 6 reports on its third column the same statistics reported for US data, but computed for the fitted values of output and the skill premium. The fitted skill premium is more volatile than output, lagging the business cycle, and the estimated contemporaneous correlation with output is positive. As the fitted value of output at time t is given by $G(e_t, k_{st}, k_{et}, u_t, s_t)$, these results show that the residual is rather important for explaining output dynamics. As is well known, much of the cyclical behavior of output is missed if one focuses solely on inputs (energy, capital, and labor) and dismisses the (Solow) residual. This pattern does not hold true for the fitted skill premium: its cyclical dynamics match well those of the *actual* skill premium. Technology

[12] These moments are computed with HP-filtered output and skill premium. Results with the other two de-trending procedures are similar and available upon request.

shocks, which greatly affect output but not the skill premium because of their neutrality, would reduce the correlation between the two and increase the volatility of output.[13]

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. This paper has clearly established the robustness of this fact. Employing three different de-trending methods (an HP filter, a band-pass filter, and deviations from an exponential trend) the correlation was found to be negative. Moreover, identifying exogenous changes in oil prices following Hoover and Perez (1992), it was found that the response of the skill premium to such a change, defined as the occurrence of a Hoover-Perez event, was negative and significant.

In addition, this paper has estimated an aggregate production function in which energy use and prices are explicitly introduced. Two key results emerge from this estimation. First, capital is more easily substituted with unskilled labor than with skilled labor. However, this finding 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, researchers have used capital-skill complementarity 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

[13] We do not have a good explanation for the lagging behavior of the fitted skill premium. Despite this behavior not being significant – the standard errors are large, the point estimate of the contemporaneous correlation (0.46) is smaller than that at one-lead (0.59). There are several factors that could be contributing to this discrepancy. Among others, abstracting from the the residential sector in our measure of fitted output (but not in the measure of actual output), or assuming that the appropriate deflator of non-residential structures is a price index of consumption goods.

observed correlation between oil prices and the skill premium: when oil prices rise, firms substitute unskilled workers for capital, and the skill premium falls.

References

[1] Berndt E. and White C. (1978), "Technology, Prices, and the Derived Demand for Energy", unpublished paper, University of British Columbia.

[2] Bernanke B., Gertler M. and Watson M. (1997), "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity* 1, pp. 91–142.

[3] Card D. and DiNardo J. (2002), "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics*, Vol. 20, no. 4, pp. 733–81.

[4] Christiano L. J. and Fitzgerald T. J. (2004), "The Band-Pass Filter," *International Economic Review*, Vol. 44(2), pp. 435-465.

[5] Clark K. and Freeman R. (1977), "Time-Series Models of the Elasticity of Demand for Labor in Manufacturing," unpublished paper, Harvard University.

[6] Denison E. (1979), *Accounting for Slower Economic Growth: The United States in the 1970s*, The Brookings Institution, Washington D.C.

[7] Dennis E. and Smith V. K. (1978), "A Neoclassical Analysis of the Demand for Real Cash Balances by Firms," *Journal of Political Economy*, 86:793–814.

[8] Edelberg W., Eichenbaum M. and Fisher J. (1999), "Understanding the Effects of a Shock to Government Purchases," *Review of Economic Dynamics* 2, pp. 166–206.

[9] Greenwood J., Hercowitz Z., and Krusell P. (1997), "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, Vol. 87(3): 342–62.

[10] Griliches Z. (1969), "Capital-Skill Complementarity," *Review of Economics and Statistics*, Vol. 51(4), pp. 465-468.

[11] Gullickson W. and Harper M. J. (1987), "Multifactor Productivity in U.S. Manufacturing, 1949-1983," *Monthly Labor Review*, pp. 18-28.

[12] Hamermesh D. (1993), *Labor Demand*, Princeton University Press, Princeton, NJ.

[13] Hoover K.D. and Perez S. J. (1994), "Post Hoc Ergo Propter once more: An Evaluation of 'Does Monetary Policy Matter?' in the Spirit of James Tobin," *Journal of Monetary Economics*, Vol. 34, pp. 47-73.

[14] Keane M. and Prasad E. (1996), "The Employment and Wage Effects of Oil Price Changes: A Sectoral Analysis", *The Review of Economics and Statistics*, Vol. 78, pp. 389-400.

[15] Krusell P., Ohanian L., Ríos-Rull J.V., and Violante G. (2000), "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, Vol. 68(5), pp. 1029-1054.

[16] Lee, K. and Ni, S. (2002), "On the Dynamic Effects of Oil Price Shocks: a Study Using Industry Level Data," *Journal of Monetary Economics*, Vol. 49, pp. 823-852.

[17] Lindquist M. (2004), "Capital-Skill Complementarity and Inequality Over the Business Cycle", *Review of Economic Dynamics*, Vol. 7, pp. 519-540.

[18] Ohanian L., Violante G. L., Krusell P., and Ríos-Rull J.V.(2000), "Simulation-Based Estimation of a Non-Linear, Latent-Factor Aggregate Production Function," in *Simulation-Based Inference in Econometrics: Methods and Applications*, Mariano R., Schuelrmann T., and Weeks M., eds. Cambridge, U.K.

[19] Polgreen L. and Silos P. (2005), "Capital-Skill Complementarity and Inequality: A Sensitivity Analysis," Federal Reserve Bank of Atlanta WP-2005-20.

[20] Polgreen L. and Silos P. (2008), "Capital Skill Complementarity and Inequality: A Sensitivity Analysis," *Review of Economic Dynamics*, Vol. 11(2), pp. 302-313. 29(4), pp. 910-929.

[21] Ramey V. and Shapiro M. (1998), "Costly Capital Reallocation and the Effects of Government Spending ," *Carnegie Rochester Conference on Public Policy*, Vol. 48, pp. 145-194. 29(4), pp. 910-929.

Table 1: Correlations: Skill Premium and Oil Prices

Filter	Entire Sample (1963-2004)	Second Subsample (1979-2004)
Exp. De-trend	-0.713 (0.066)	-0.690 (0.091)
BP De-trend	-0.434 (0.135)	-0.397 (0.189)
HP De-trend	-0.312 (0.154)	-0.343 (0.173)

Notes: Correlation coefficient between oil prices and the skill premium for two different sub-samples (across columns) and three different de-trending procedures (across rows). The HP-filter uses a smoothing parameter of 100. The bandpass filter eliminates fluctuations occurring at periods shorter than three years or longer than 35.

Table 2: 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)

Notes: Prior distributions for the parameters of the structural model. The indicator variable $\chi_A(x)$ takes the value one if the random variable x belongs to set A, and zero otherwise.

Table 3: 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

Notes: The first column gives posterior means of σ and ρ, and the OLS estimate of ν using relationship (11). The second column gives posterior standard deviations for σ and ρ, and the standard error (s.e.) for the OLS estimate of ν.

Table 4: Correlations: Oil Prices vs. Fitted Skill Premium

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)

Notes: The first column gives the mean correlation between actual oil prices and the fitted skill premium resulting from averaging across correlations computed for all draws of the model's parameter vector. We de-trend oil prices and the fitted skill premium using a different procedure in each of the three rows. The second column computes the fitted skill premium once using the mean of the estimated parameters and computes its correlation with actual oil prices. Standard errors in this case are computed using GMM.

Table 5: Relative Volatility ($\frac{\sigma_p}{\sigma_{SP}}$): Oil Prices vs. Fitted Skill Premium

De-trending Proc.	Posterior Distribution	GMM s.e.'s
Exp. De-trending	4.354 (0.815)	4.109 (0.397)
BP Filter	8.561 (0.546)	8.541 (1.100)
HP Filter	7.997 (0.488)	7.977 (1.325)

Notes: The first column gives the mean ratio of volatilities of actual oil prices and the fitted skill premium resulting from averaging across ratios of volatilities computed using all draws of the model's parameter vector. We de-trend oil prices and the fitted skill premium using a different procedure in each of the three rows. The second column computes the fitted skill premium once using the mean of the estimated parameters and computes its volatility relative to that of oil prices. Standard errors in this case are computed using GMM.

Table 6: Output vs. Skill Premium (U.S. Data and Fitted Values)

	U.S. Data	Model
σ_{GDP}/σ_{SP}	1.159 (0.155)	0.644 (0.096)
$Corr(SP_t, GDP_t)$	0.206 (0.146)	0.460 (0.148)
$Corr(SP_{t-1}, GDP_t)$	0.446 (0.105)	0.291 (0.131)
$Corr(SP_{t+1}, GDP_t)$	-0.106 (0.175)	0.590 (0.079)

Notes: The first column displays the ratio of the standard deviation of US output and US skill premium, the contemporaneous correlation between those two variables, and the correlation of the US skill premium with one lead and one lag of US output. Standard errors computed by GMM in parentheses. The second column gives the analogous moments using the fitted skill premium and the fitted output which were computed with the mean values of the estimated parameters.

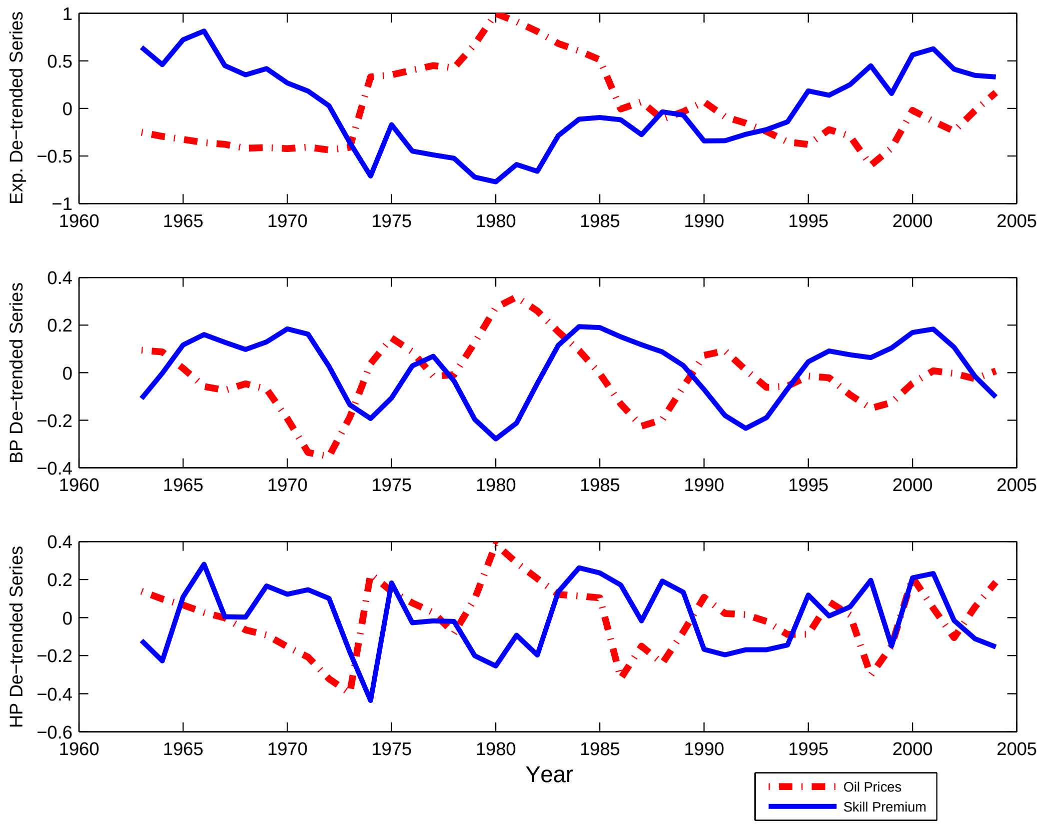

Figure 1: De-trended energy prices (dashed and dotted line) and skill premium (solid line). Three different de-trending methods: exponential de-trending (top panel), band-pass filtering (medium panel), and HP filtering (bottom panel)

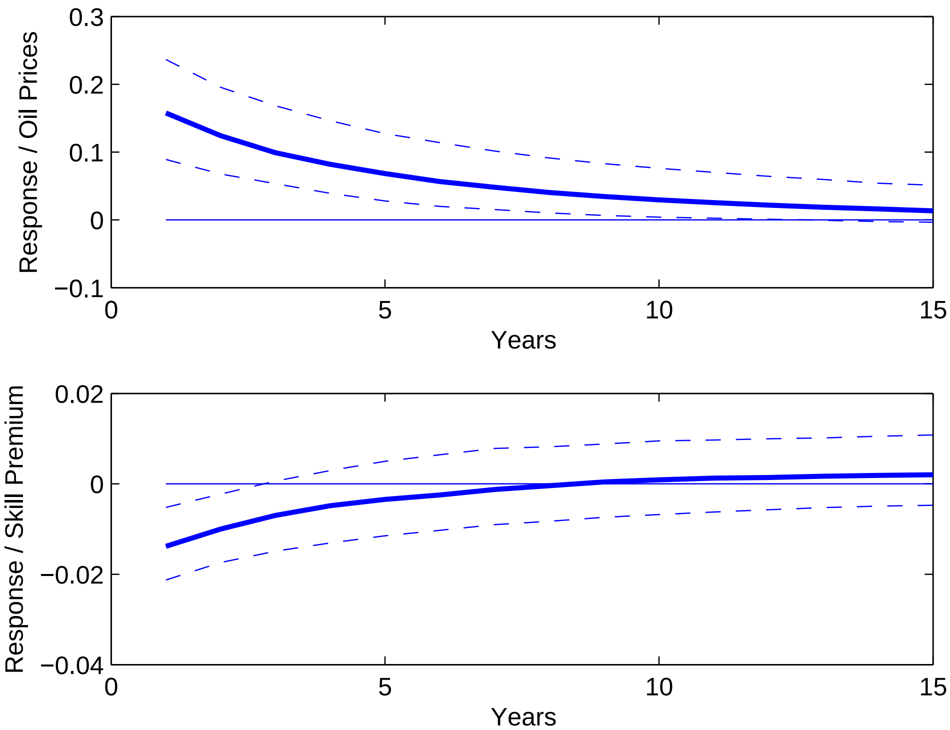

Figure 2: Responses of oil prices – top panel – and the skill premium – bottom panel – to the onset of a Hoover-Perez episode over a 15-year horizon. The solid line is the median response and the two dotted lines represent 66% confidence bands

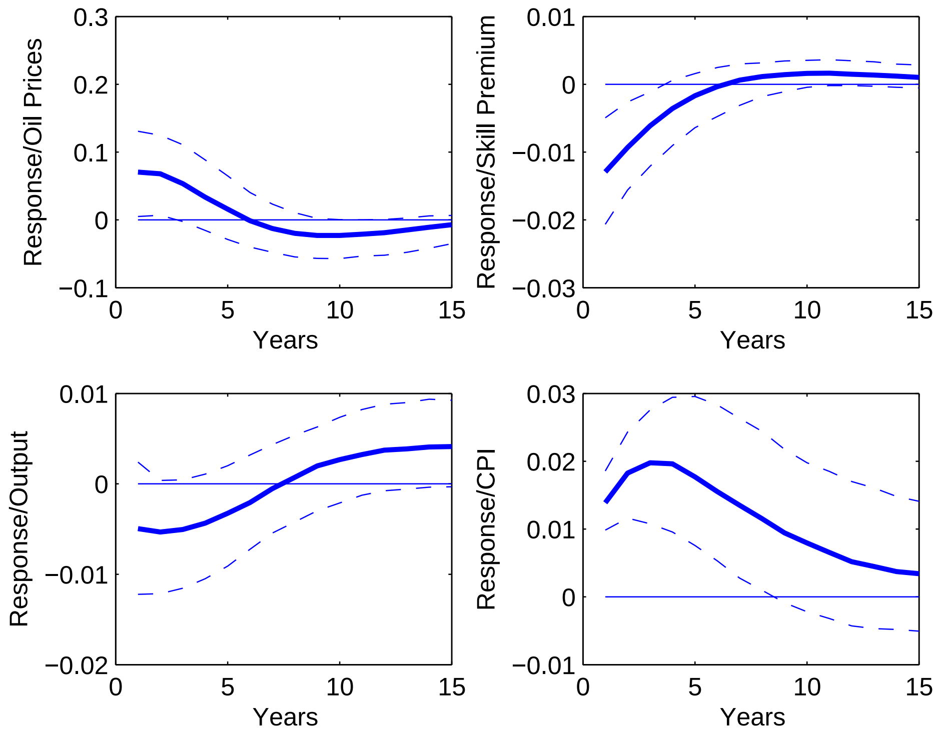

Figure 3: Response of oil prices, skill premium, real output, and the consumer price index (left to right, top to bottom order), to the onset of a Hoover-Perez episode. The solid line is the median response and the two dotted lines represent 66% confidence bands.

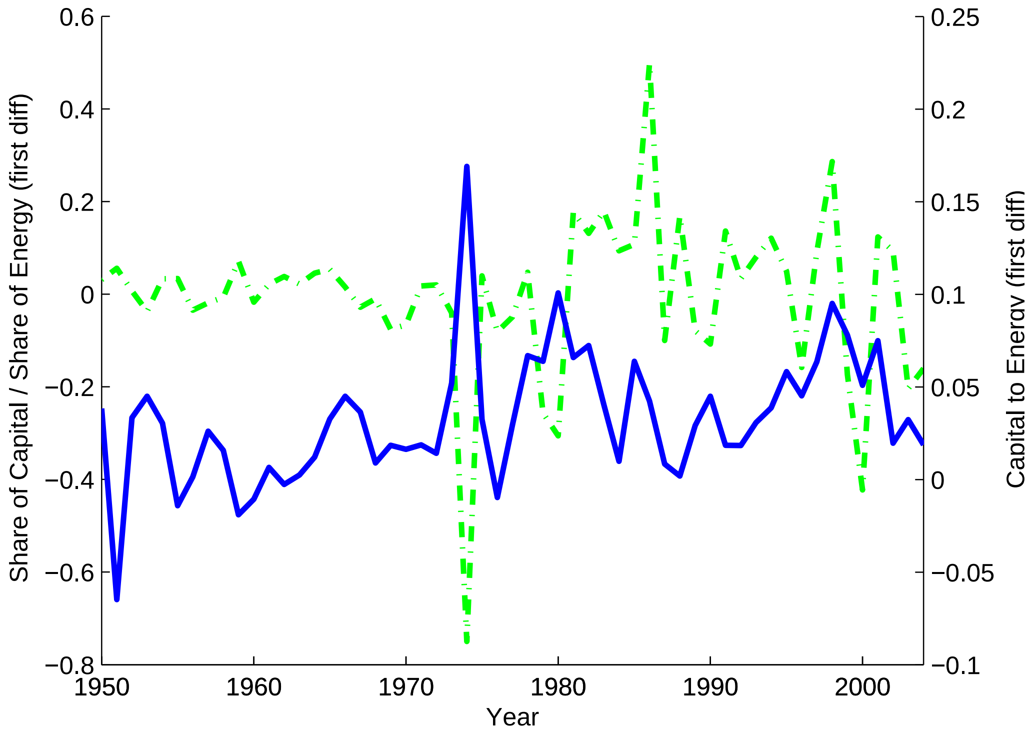

Figure 4: Two series to estimate ν. The dashed-dotted line (left-hand axis) is the ratio of the share of capital in output to the share of energy in output. The solid line (right-hand axis) is the ratio of capital to energy.