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

Linnea Polgreen*†

University of Iowa

Pedro Silos‡

Federal Reserve Bank of Atlanta

August 18, 2006

Abstract

Higher oil price shocks benefit unskilled workers relative to skilled workers. This is reflected by the strong, negative correlation between energy prices and the skill premium at the business cycle frequency. This correlation is robust to different detrending procedures. We construct and estimate a model economy with heterogeneous skills and energy use and study its business cycle implications, in particular the cyclical behavior of oil prices and the skill premium. In our economy, the skill premium and the ratio of hours worked by skilled workers to hours worked by unskilled workers are both negatively correlated with oil prices over the business cycle. The key ingredient for the skill premium and energy prices to move in opposite directions is the larger substitutability of capital for unskilled labor than for skilled labor. With our estimates, even when energy and capital are fairly good substitutes, the negative correlation arises.

Keywords: skill-heterogeneity, energy prices, economic fluctuations, skill premium

JEL Classification: E24, E32, J24

---

*Thanks to Steve Durlauf, Karsten Jeske, B. Ravikumar, Víctor Ríos-Rull, Ellis Tallman, Robert Tamura, and seminar participants at the Midwest Macroeconomic Meetings and the Econometric Society Summer Meetings. from whom we have received many useful comments.

†Department of Economics, University of Iowa, Iowa City, IA 52245, tel: 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

The average growth rates for the skill premium and oil prices have been positive for the last forty years. However, when examined closely, oil prices and the skill premium move in opposite directions, displaying a very strong, negative correlation. This negative correlation is found specifically at the business cycle frequency, and it is robust to different detrending methods.

We employ a version of the stochastic growth model to examine and quantify the mechanism by which the negative correlation arises, and we conclude that capital-skill complementarity drives the dynamics of oil prices and the skill premium. The model features a technology that uses energy as an input, oil price variations, heterogeneity in skill and a larger complementarity between capital and skilled labor than between capital and unskilled labor. In our model, oil prices affect our energy-capital composite[1], which in turn is the driving force behind changes in the skill premium. Specifically, as oil prices increase, energy use decreases. As long as capital and energy are not too substitutable[2], oil price increases will reduce the capital-energy composite. Because unskilled labor is more easily substituted for the capital-energy composite than skilled labor[3], the demand for unskilled labor relative to skilled labor increases, raising the relative wage of the unskilled and decreasing the skill premium.

To analyze the behavior of the skill premium and oil prices we use a general equilibrium model instead of focusing only on the production side of the economy. The reason we use this framework is that labor hours, for reasons outlined below, are important to consider: our model is consistent with not only with the joint behavior of the skill premium and oil prices, but also with the relationship between hours worked by skill group and energy

[1]We combine capital and energy into one term in order to facilitate parameterization. This is discussed in section 3.

[2]There is mixed evidence regarding the complementarity of energy and capital (e.g. Pindyck (1979) suggests that in the short run, energy and capital are substitutes, but in the long run they are complements), so we examine cases for which the elasticity of substitution between capital and energy is positive, zero and negative.

[3]See Hamermesh (1993) for a review of the literature.

prices observed during most of the sample.

The model presented here is able to account for the negative correlation between cyclical movements in oil prices and the skill premium. Interestingly in our model, although the skill premium is mainly driven by movements in energy prices, movements in output are almost entirely the result of technology shocks, in agreement with the basic premise of real business cycle models. The model is also consistent with more standard business cycle observations, such as the relative volatilities of different aggregates, as well as their cross-correlations with output. We also contribute to the literature on income inequality over the business cycle (e.g. Castañeda, Díaz-Giménez and Ríos-Rull(2003)). We find that although the ratio of energy expenditures to total capital in the US economy is small, the variability of oil prices is large relative, for instance, to variations in the Solow residuals. As already mentioned, the volatility of the skill premium is driven largely by the reaction of the capital-energy composite to changes in oil prices. We find that the volatility of energy expenditures is a large determinant in the overall variability of the skill premium; thus, the results from our model more closely resemble the data than previous research.

Energy prices have been largely ignored in the study of the skill premium. To our knowledge this is the first paper that examines the relationship between cyclical movements in the skill premium and oil prices within an equilibrium model of economic fluctuations. Nevertheless, our work is related to previous studies that have analyzed the role of energy in real business cycle economies (e.g. Kim and Loungani (1992)) or have analyzed the behavior of the skill premium in equilibrium models. Here the two most notable examples are Krussell et al. (2000) and Lindquist (2004)[4]. There is also related empirical work using panel data. In opposition to our findings, Keane and Prasad (1996) find that skilled, rather than unskilled workers, gain during oil price increases. Keane

[4]Other related work include Prasad(1996), who analyzed the implications of skill heterogeneity in a business cycle model for the cyclical behavior of in productivity and the real wage, and Castro and Coen-Pirani (2005) who undertake a careful evaluation of the change in the cyclical behavior of aggregate skilled hours after 1984.

and Prasad use data covering the period 1966-1981, and we believe that their results are driven by the time period considered. This will be discussed in more detail in the next section.

2 Energy Prices and the Labor Market

The focus here is to summarize the behavior of the skill premium and energy prices over the last decades. The skill premium is a weighted ratio of skilled wages to unskilled wages[5]. 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), 1963–2001.

Data on energy prices and usage come from the US Government Energy Information Administration. We were able to obtain annual data from 1949 to 2001 for prices and quantities of oil, coal and natural gas. The reason to focus only on fossil fuels is that they represent almost 85% of overall energy consumption in the US. In addition, we will be able to compare our results to previous studies, such as Kim and Loungani (1992). The price of energy used throughout the analysis is a Laspeyres index of the prices of those three main energy sources. 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. All results discussed would hold by using oil prices instead of the measure used here. The final energy price index used to map model and data was the result of dividing the constructed energy price index by the Gross Domestic Product deflator.

Deviations of the energy price index relative to its HP-trend are shown in Figure 1. 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.

[5]Details are provided in the appendix.

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%). Overall, energy prices have displayed a large amount of volatility over the last three decades.

Figures 2 - 4 show the detrended skill premium and energy prices, using three types of de-trending methods: deviations from an exponential trend, a (log)HP-filtered series and a (log)band-pass-filtered series. Correlations are negative and in some cases surprisingly strong. For instance, the correlation between the skill premium and energy prices is -0.77 when measured as deviations from an exponential trend. With the other two methods correlations are not as strong but still significant and on the order of -0.4.

As mentioned in Keane and Prasad (1996), the negative correlation between the skill premium and energy prices could be an artifact of aggregation. Wages are only available for the employed, so our skill premium compares the wages of the skilled who are employed to the unskilled who are employed. When energy prices rise, firms that need to cut costs may lay off the lowest-skilled and lowest-paid employees, raising the average wage of the unskilled. The skill premium would rise, even if wages have not changed at all. However, support for our argument is found by examining the labor input ratio, defined as the hours worked by a skilled worker divided by the hours worked by an unskilled worker. Figure 5 shows the detrended labor input ratio and the detrended energy prices. The correlation between those two series is -0.1, which basically implies that the hours ratio and oil prices are uncorrelated. For most of the series, however, the labor input ratio and oil prices appear to be negatively correlated. The first oil shock is an exception: in 1974 the labor input ratio increased as prices increased. The first oil shock is the primary cause of the weak correlation. We do not have an explanation for this observation, but firms might have perceived that the oil price shock was temporary. It is usually more difficult to replace high-skilled workers than low-skilled ones, so it might have been optimal to lay off relatively more of the unskilled, increasing the labor input ratio. It is clear from the figure that the labor input ratio increased (decreased) when oil prices decreased (increased), approximately after 1978. The importance of the first oil shock in explaining the weak correlation is more clear when we compute a sequence of 9-year rolling correlations with the first ending in 1972 and the last one ending in 2001. Both the raw data and a spline-smoothed approximation of these correlations are plotted in Figure 6. After the first oil price shock there is a sharp decline in the correlation between the labor ratio and oil prices, and that correlation stays below zero until the end of the sample. In fact, using data from 1978 until the end of the sample, the correlation is -0.41. This anomalous behavior during the first oil price shock could partly explain Keane and Prasad's results. Their NLSY data covered only the period 1966-1981.

Table 1: Volatilities and Correlations with Output
Annual Data (1963-2001)

Variable	Std. Dev. rel. GDP	Correl. with GDP
Consumption	0.59	0.86
Investment	2.96	0.93
Unskilled Hours	0.34	0.73
Skilled Hours	0.27	0.60
Energy Use	1.06	0.31
Energy Prices	8.91	-0.40
Skill Premium	0.82	0.19
Hours Ratio	0.20	-0.36

Finally, Table 1 reports some business cycle statistics for several macroeconomic variables. Consumption (defined as non-durables and services expenditures), Fixed Investment, Output (the sum of Consumption and Investment) and Energy Use were transformed into per-capita quantities by dividing by the US population, deflated using the GDP deflator, logged and detrended using an HP filter.

Consumption, investment, hours and energy use are all procyclical, although the correlation of energy use with contemporaneous output is rather weak (0.31). The skill premium could be considered almost acyclical with an even weaker correlation with GDP of 0.19. This fact is consistent with other studies of the skill premium over the business cycle, such as Lindquist (2004) who finds, with quarterly data, a correlation closer to zero.

Higher energy prices are associated with recessions, and this is reflected in the negative correlation between oil prices and output (-0.30). Finally, the hours ratio (skilled hours over unskilled hours) is also counter-cyclical with a correlation with contemporaneous output of -0.36.

Regarding the relative volatilities, consumption and hours are less volatile than output and investment. The table shows that energy use is roughly as volatile as GDP. However, energy prices are exceptionally volatile: their volatility relative to output's is almost 9.

3 The Model

The economy is populated by a continuum of infinitely-lived agents which are of two types: skilled and unskilled. Within each type, all agents are identical and individuals may not transit across types. Denote by s the fraction of skilled and by u the fraction of unskilled, with $s+u=1$. Agents value consumption and leisure. They rank their options according to the utility function $u(c_{t,j}, 1-h_{t,j})$, where $c_{t,j}$ and $h_{t,j}$ represent consumption and time spent at work respectively for an agent of type j, $j \in \{u, s\}$. Agents are endowed with one unit of time each period, which they divide between work and leisure, and both types discount the future with a factor β.

There is a representative firm that produces output (Y) using energy (E), capital (K), skilled hours (H_s) and unskilled hours (H_u). Technology is represented by the following constant returns to scale production function:

$$Y_t = z_t G(K_t, E_t, H_{s,t}, H_{u,t}) \tag{1}$$

In the above expression z_t is a random variable representing neutral technological change. The firm uses aggregate hours as their input and therefore $H_{j,t} = jh_{j,t}$, for $j \in \{s, u\}$. We deviate slightly from previous studies of the skill premium by aggregating all types of capital into one variable K_t [6].

[6] We are aware of the advantages of separating total capital into structures and equipment; the faster decline of equipment prices helps to understand the evolution of the skill premium at lower frequencies. We have chosen this simpler approach because we believe that for the goal of this paper, it suffices to have only one type of capital. Also, it clarifies the exposition.

In this economy, markets for goods and factors are competitive. We do not explicitly model the underlying production of energy, and we assume that it involves forgoing a certain amount of consumption and physical capital. The amount needed, however, varies because the relative price of energy (p) with respect to the consumption good evolves exogenously.

The absence of distorting taxes, externalities, etc. allows us to invoke the welfare theorems and solve the associated social planner's problem. The planner maximizes the weighted sum of utilities for the two types of agents by choosing sequences of capital, consumption, labor and energy. Formally the problem can be stated as:

$$\max_{\{c_{t,s},c_{t,u},h_{t,s},h_{t,u},E_t,K_{t+1}\}} E\sum_{t=0}^{\infty}\beta^t\{\Psi s\left[u(c_{t,s},1-h_{t,s})\right]\}+(1-\Psi)u\left[u(c_{t,u},1-h_{t,u})\right]\} \quad (2)$$

s.t.

$$sc_{t,s}+uc_{t,u}+p_tE_t+K_{t+1}\leq Y_t+(1-\gamma)K_t$$

$$H_{t,j}=jh_{t,j}$$

$$Y_t=z_tG(K_t,E_t,H_{t,u},H_{t,u})$$

$$H_{t,s},H_{t,u},C_{t,s},C_{t,u}>0$$

Denoting by η the vector of exogenous shocks $(\log z_t,\log p_t)'$, we assume that it follows a first-order Markov process:

$$\eta_t=\Phi\eta_{t-1}+\nu_t \quad \nu_t\sim N(0,\Omega). \quad (3)$$

Innovations to technology and oil prices can be contemporaneously correlated, i.e. Ω is unrestricted. The companion matrix Φ is restricted to be diagonal for simplicity. An equilibrium for this model is a set of decision rules for the endogenous variables, given exogenous shocks and parameters, which solve the planner's problem, and a set of factor prices that are equal to the marginal products of skilled labor, unskilled labor and capital.

4 Parameterization

We restrict preferences to be of the logarithmic class with separability between consumption and leisure,

$$u(c_{t,j}, 1 - h_{t,j}) = \theta log(c_{t,j}) + (1 - \theta) log(1 - h_{t,j}) \quad i \in \{j, s\}$$

with the parameter θ representing the "expenditure" share of each of the two goods. Note that preferences are identical for each of the two types of agents.

Output is obtained using capital, energy and labor and produced according to the following nested-CES production function,

$$Y_t = z_t\{\xi(\alpha\tilde{K}_t^{\phi} + (1 - \alpha)H_{t,s}^{\phi})^{\frac{\delta}{\phi}} + (1 - \xi)H_{t,u}^{\delta}\}^{\frac{1}{\delta}}$$

The variable $\tilde{K}_t$ is the capital-energy composite:

$$\tilde{K}_t = \{\mu K_t^{\nu} + (1 - \mu)E_t^{\nu}\}^{\frac{1}{\nu}}$$

We write production in this way because we will (rather loosely) interpret the capital-energy composite used here as the measure of capital used in other studies in order to assign values to some parameters.

Obtaining quantitative conclusions requires parameterizing the model in a realistic way. In principle, it is possible to estimate the model by maximum likelihood using well-known methods. However, in early attempts, the likelihood function turned out to be ill-behaved and we attacked the problem using a simulated quasi-maximum likelihood method, augmenting the estimation with prior distributions over the structural parameters of the model. We provide a brief description of the method below and interested readers are referred to Smith's (1993) work for a more detailed explanation. The estimation procedure, in short, maximizes a likelihood function (the quasi-likelihood function) that differs from the exact likelihood of the model economy. Let θ be a vector of structural parameters (describing preferences, technology, etc,...) and let $\{\tilde{y}\}_{t=1}^{S} = f(\theta, \{\eta_t\}_{t=1}^{S})$ be the output from the model, a vector of time series of GDP, employment, energy use, etc,... of length S, which is an unknown function of the structural parameters and a sequence of realizations of the two shocks. The estimation procedure fits a reduced-form statistical model to $\{\tilde{y}\}_{t=1}^{S}$, in our case a VAR, with a well defined likelihood function yielding a set of parameters $\beta(\theta)$ (in our case, the OLS estimates of the VAR). Denote this likelihood function by $L(\{\tilde{y}\}_{t=1}^{S}; \beta(\theta))$. The quasi-likelihood function of the model is $L(\{Y\}_{t=1}^{T}; \beta(\theta))$, where $\{Y\}_{t=1}^{T}$ is the empirical counterpart of $\tilde{y}$ obtained from actual US data.

We augmented this quasi-likelihood function with prior distributions over the structural parameters, $p(\theta)$. We believe that incorporating prior information about the parameters is an advantage, not a drawback, of this Bayesian approach and we summarize this information in the form of probability density functions. Coupling the quasi-likelihood function with the prior distributions we obtained our "quasi-posterior" distribution $P(\theta|\{Y\}_{t=1}^{T}) \propto L(\{Y\}_{t=1}^{T}|\theta)p(\theta)$. We simulated a long sequence of draws from the quasi-posterior distribution using well-known sampling procedures (see, for example, Fernández-Villaverde and Rubio-Ramírez (2004))[7].

The entire vector to be estimated was $(\beta, \psi, \theta, \xi, \alpha, \mu, \nu, \delta, \gamma, \phi, \rho_p, \rho_z, \sigma_{pz}, \sigma_p, \sigma_z)'$. To facilitate the estimation we fixed some parameters that have a clear empirical counterpart and whose values have been estimated elsewhere in the literature. For instance, the depreciation rate of capital γ was set at 0.1, and the discount factor β was set at 0.96, widely used values when the frequency in a model is annual. The three parameters that drive the elasticities of substitution between the different factors are ϕ, δ and ν. We loosely assigned ϕ, the parameter driving the elasticity of substitution between skilled labor and the capital-energy composite, a value consistent with estimates found in previous studies. Krusell et al. (2000) estimate ϕ to be -0.45, while Polgreen and Silos (2006), building on Krusell et al.'s analysis, find values for ϕ between -0.16 and -0.60. We used the Krusell

[7]For a given vector of parameter values, the model's solution is found by log-linearizing the optimality conditions and the constraints, and solving for the expectation functions using the methods described in Klein (2000). A Technical Appendix at the end describes the precise equations used when solving the model.

et al. estimate. Similarly, δ, the elasticity of substitution between unskilled labor and capital, was assigned a value of 0.5. The literature has not agreed on how substitutable energy and capital are. Work using time-series data find values that imply that both inputs are more complementary than a Cobb-Douglas would imply, while conclusions from cross-sectional studies point towards more substitutability than Cobb-Douglas. We have set the parameter ν to -0.7, the estimated value in Morrison and Berndt (1981), also used in Kim and Loungani (1992). However, the dependence of the quantitative conclusions on the value of ν will be the object of a detailed discussion below. Finally the fraction of skilled workers s was set to 0.28, the average for our sample period. We also fixed the persistence parameter and the variance of the noise in the oil price shock equation (3): ρ_p was fixed at 0.846 and σ_p^2 was set at 0.062.

The remaining parameters $(\psi, \theta, \xi, \alpha, \mu, \rho_z, \sigma_{pz}, \sigma_z)'$ were estimated. We used three different observables: real output, consumption and energy use [8]. The data were logged and HP-filtered prior to estimation.

The prior distributions for the parameters were all independent normal and gamma. We did perform some prior predictive analysis to guide us in the choice and shape of $p(\theta)$. As a result, we centered the distribution for μ, the weight of capital in the capital-energy composite, at 0.97 with a standard deviation of 0.05, in order to attain a small energy to capital ratio, as observed in US data. We had less prior information about the remaining weights in the production function: α and ξ had normal prior distributions centered at 0.5 with a standard deviation of 0.1; the same distribution as the planner's weight Ψ. The parameter controlling preferences towards consumption, θ was given a prior mean of 0.7 with a standard deviation of 0.05. Regarding eq. (3), the prior mean of ρ_z was 0.9 and the standard deviation 0.03; the covariance between the two shocks (σ_{pz}) was endowed with a normal distribution with a mean of -0.001 and a standard deviation of 2×10^{-4}. Finally, the prior distribution for the variance of the productivity shock was gamma with parameters 10 and 1×10^{-4} which implies a mean of 1×10^{-3} and a standard deviation

[8] A section below provides some sensitivity analyses to changes in the choice of observables.

of 3×10^{-4}. All the normal distributions were truncated to the appropriate regions.

The posterior means and standard deviations of the estimated parameters are given in Table 2. Although the prior distributions were quite informative, for several parameters the data centered the posterior distributions far away from the prior means. For example, the weight of capital μ has a posterior mean of 0.90 which is more than two standard deviations away from the prior mean. The distributions of the other production function parameters, ξ and α were each displaced by approximately one standard deviation. The other parameters' distributions had shifts of smaller magnitudes.

Table 2: Posterior Means and Standard Deviations

Parameter	Mean (Std. Dev.)
Ψ	0.472 (0.074)
ξ	0.655 (0.102)
α	0.611 (0.046)
μ	0.903 (0.013)
θ	0.749 (0.049)
ρ_z	0.910 (0.013)
σ_{pz}	-1.2×10^{-3} (2×10^{-4})
σ_z^2	1.2×10^{-4} (2.0×10^{-5})

5 Results

The quantitative evaluation of the model was done in a standard way. After solving for the policy functions and simulating the shock processes using the parameters presented above, we obtained a set of time series of interest. We treat these series in levels the same way as the true data: we first logged them and then HP-filtered them. We present moments (standard deviations and contemporaneous correlations with output[9]) of the deviations of variables from their HP trend. Table 3 presents measures of volatility for a few macroeconomic aggregates, with a standard error in parentheses [10]:

[9]We have decided not to overwhelm the reader with columns of data on cross-correlations with GDP at different leads and lags. These are, of course, available upon request for any of the parameterizations in the paper.

[10]The standard errors are posterior standard deviations of the standard deviations themselves. For each of the draws of θ, the vector of structural parameters, we solved for the decision rules and simulated

Table 3: Standard Deviations (in %)

Variable	Std. Dev. (Std. Error)
Y	5.35 (0.52)
C	1.83 (0.24)
I	18.92 (2.21)
H_s	0.87 (0.27)
H_u	0.53 (0.23)
E	18.26 (0.32)
w_s/w_u	2.75 (0.37)
H_s/H_u	0.29 (0.21)

The model's implications regarding standard deviations are broadly consistent with US data: consumption and hours are less volatile than income, which is less volatile than investment. Quantitatively, energy use is too volatile in the model with a standard deviation relative to that of GDP of more than three. In the data their standard deviations are about the same. Finally, both the skill premium and the hours ratio are significantly less volatile than output, as is observed empirically.

The contemporaneous correlations with output, these are shown in Table 4

Table 4: Correlations with Output

Variable	Correlation (Std. Error)
C	0.708 (0.044)
I	0.819 (0.043)
H_s	0.826 (0.054)
H_u	0.428 (0.113)
E	0.783 (0.044)
w_s/w_u	0.755 (0.082)
H_s/H_u	0.587 (0.160)

Consumption, investment and production inputs are all procyclical as in the data, although the correlation between energy use and output is much stronger in the model.

the economy, therefore obtaining an entire distribution of the standard deviations and correlations with output of any aggregate variable. The standard errors are not a result of simulating time series of different lengths and then averaging over them, as it is sometimes done in the macroeconomics literature.

The same can be said for the skill premium, which is mildly procyclical in the data but has a correlation higher than 0.7 in the model. Finally, it is the correlation between output and the hours ratio in which the model fares worst: it is countercyclical in the data but procyclical in the model. A subsection below will analyze whether increasing the substitutability between capital and energy can improve along these dimensions.

The correlations between the skill premium and energy prices, as well as the relationship between oil prices and the relative labor ratio, are presented in Figure 7. There is a great deal of uncertainty in the model about the value of the correlation between the hours ratio and oil prices with the posterior density for this correlation covering a wide range from -1 to 0, and with a substantial amount of mass between -0.1 and 0. The posterior distribution for the correlation between the skill premium and oil prices is much tighter, and according to the model values larger than -0.8 are unlikely.

Finally, our model has also implications for the volatility of the skill premium itself. Table 1 reports the skill premium's volatility relative to GDP in annual US data of 0.82. As is clear from Table 3, the model delivers a volatility substantially lower than what is observed in the data: the ratio of the volatility of the skill premium to that of GDP is only 0.17.

The model contains both TFP and oil-price shocks, but we want to determine how much of the variance of the skill premium is attributable to oil-price shocks only. Results for this experiment are presented in Table 6. We show the standard deviation of the skill premium relative to that of output and the standard deviation of the skill premium itself. These are computed for three distinct economies: one where the only shocks are oil price shocks, one where the only shocks are TFP shocks and one where both shocks hit the economy.

Table 5: Relative Contributions of Shocks to Skill Premium Variation

Shock	$std(SP)/STD(GDP)$	$std(SP)$ (in %)
p	0.83	0.85
z	0.07	0.33
Both	0.17	0.86

Energy price shocks are an important source of fluctuations in the skill premium. Quantitatively they are considerably more important than TFP shocks. In fact, in economies where the only shocks are energy shocks, the volatility of the skill premium relative to that of GDP matches the data; we observe a ratio of approximately 0.80. Even in the case of high substitutability between capital and energy, this ratio is only 0.22 when both shocks are present. Because of the neutrality of TFP shock, one would expect that energy prices would be relatively more important in explaining movements in the skill premium. Nevertheless, quantitatively the difference is large.

5.1 Alternative Parameterizations

In this section we focus on analyzing the effects of changing two elements from the previous parameterization. First, we will explore the implications of reducing the complementarity between capital and energy. Second, we will provide alternative estimates for the structural parameters of the model as a result of using different time series in the estimation procedure.

Let us start by assessing how the correlation between the skill premium and oil prices and the correlation between the hours ratio and oil prices change when energy and capital are more substitutable than a value of ν equal to -0.7 implies. Figure 8 shows the posterior distributions of the correlation between the skill premium and oil prices for three different values of ν: -0.7,0, and 0.3 [11]. These three values correspond, respectively, to more, about the same and less complementarity than a Cobb-Douglas capital-energy aggregator. The remaining parameters in the model were fixed at the values shown in Table 2.

There are two features worth noticing about the graph. The first is the tightness of the posterior distribution for the higher complementarity cases. The second is the non-

[11] The value of ν was not exactly zero. We set it at 0.001.

monotonicity in the average values of the correlations. The "close-to-Cobb-Douglas" case implies a smaller correlation than the higher complementarity case, which in turn implies a smaller correlation than the higher substitutability case. In understanding the relationship between the skill premium and energy prices it is key to understand the behavior of the capital-energy composite in the face of an oil price shock. More substitution implies that when the economy is hit by a price shock, capital and energy will move in opposite directions. Complementarity implies that they will move together. By keeping all other parameters fixed, the larger the substitutability between capital and energy, the smaller the volatility of the capital-energy composite, and therefore the smaller the volatility of the skill premium. When we change ν from -0.7 to 0, the decrease in volatility is large relative to the increase in the covariance. The latter is still negative and the small standard deviation results in a strong negative correlation. In the higher substitution case the decrease in the covariance dominates the fall in the standard deviation yielding a larger correlation.

Figure 9 shows the posterior distributions for the correlation between the hours ratio $(\frac{H_s}{H_u})$ and oil prices. The uncertainty is large and all cases have at least some mass at values close to zero. The increasing correlation when we increase substitutability, which is an expected result, is clear from the picture. Concluding, we can say that the assumption of high capital-energy complementarity is not necessary to generate a negative correlation between oil prices and the skill premium.

Let us now turn to the sensitivity of the model's results to changing the time series used in the estimation. Recall that our choice in the previous section was real per capita consumption, output and energy expenditures. The two panels in Table 6 show the standard deviations and the contemporaneous correlations with output for the same aggregates as those shown in Tables 3 and 4. Each column uses different combinations of time series in the estimation procedure described above. The series are real per capita consumption (C), investment (I), output (Y), energy expenditures (E) and employment (N). We produced results using output, consumption and employment (YCN); output, investment and energy expenditures (YIE); and output, investment and employment (YIN).

Table 6a: Standard Deviations (alt. estimation)

YCN: Output, Consumption, Employment

YIE: Output, Investment, Energy

YIN: Output, Investment, Employment

Variable	YCN	YIE	YIN
Y	7.26 (0.88)	4.97 (0.67)	3.78 (0.27)
C	2.06 (0.16)	1.28 (0.12)	1.72 (0.13)
I	16.9 (1.10)	26.7 (2.04)	9.91 (0.79)
H_s	0.88 (0.14)	2.09 (0.40)	0.18 (0.04)
H_u	0.99 (0.17)	0.21 (0.06)	1.22 (0.14)
E	18.3 (0.25)	17.8 (0.27)	16.0 (0.19)
w_s/w_u	2.23 (0.25)	1.46 (0.35)	1.89 (0.18)
H_s/H_u	1.05 (0.11)	2.00 (0.38)	1.17 (0.14)

Table 6b: Correlation with Output (alt. estimation)

Variable	YCN	YIE	YIN
C	0.74 (0.02)	0.76 (0.05)	0.81 (0.02)
I	0.70 (0.04)	0.66 (0.03)	0.90 (0.03)
H_s	0.68 (0.05)	0.61 (0.06)	0.90 (0.02)
H_u	0.40 (0.07)	0.56 (0.03)	0.23 (0.11))
E	0.59 (0.06)	0.54 (0.07))	0.84 (0.06))
w_s/w_u	0.60 (0.07)	0.30 (0.12)	0.90 (0.04)
H_s/H_u	0.19 (0.14)	0.58 (0.06)	-0.11 (0.13)

Most of the qualitative results shown in Tables 3 and 4 still hold. In particular, consumption, energy use, investment and employment are procyclical. Investment and energy use are substantially more volatile than the other aggregates, although there are some large quantitative differences. For instance, the ratio investment to output volatility is 5.4 for the YIE parameterization, substantially larger than for the other cases. On the other hand, the standard deviation of skilled hours relative to unskilled hours seems to be quite sensitive to the series used: a factor of 10 is almost reversed when going from the third to the fourth column in Table 6a. Skilled hours are even more volatile than consumption in the YIE case. Interestingly, in one of the dimensions along which the model had the most problems matching the data the point estimate has now the correct sign. The YIN parameterization implies a negative correlation of the relative hours ratio and output, however, the standard deviation of this estimate is large and zero is within one standard deviation of its mean.

Finally, Figures 10 and 11 show the posterior distributions of the correlations of the skill premium with oil prices and the correlation of the relative hours ration with oil prices, respectively. All the parameterizations show negative correlations between the skill premium and oil prices and distributions differ only on their tightness. The sign of the correlation coefficient between oil prices and the relative hours ration is more sensitive to alternative parameterizations: the YIN case displays a distribution for that correlation that has very small mass for negative values. The other two cases show the opposite, with one of the distributions (YCN) having a very small variance centered around -0.9.

6 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, approximately, the business cycle frequency. This fact is robust to different de-trending methods, and this correlation is surprisingly strong.

Previous researchers using different data sets and sample periods have found the opposite: a negative oil price shock benefits the skilled worker relative to the unskilled worker. However, we have shown that the negative correlation is what a reasonably parameterized real business cycle model predicts. While we could find parameterizations in which there exists a positive correlation between the skill premium and energy prices, when the model is forced to match a few moments from the data, the negative correlation obtains, even for a high degree of substitutability between energy and capital. A key ingredient in the model is the larger substitutability of capital for unskilled labor than for 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)).

Finally, we have shown the importance of energy-price shocks on explaining the variation of the skill premium. Energy-price shocks are quantitatively much more important than TFP shocks, and in fact, economies where the only source of uncertainty are oil price shocks, we are able to account for the volatility of the skill premium relative to that of GDP.

A Data

The skill premium is calculated using a method from Polgreen and Silos (2006). 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, age^2, 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 wages[12].

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 series[13], N.

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

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.

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

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

[13]The 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.

B Model Solution

The planner solves the following problem:

$$\max_{\{c_{t,s},c_{t,u},h_{t,s},h_{t,u},E_t,K_{t+1}\}} E\sum_{t=0}^{\infty}\beta^t\{\Psi s\left[u(c_{t,s},1-h_{t,s})\right]\}+(1-\Psi)u\left[u(c_{t,u},1-h_{t,u})\right]\}$$

s.t.

$$sc_{t,s}+uc_{t,u}+p_tE_t+K_{t+1}\leq Y_t+(1-\gamma)K_t$$

$$H_{t,j}=jh_{t,j}$$

$$Y_t=z_tG(K_t,E_t,H_{t,u},H_{t,u})$$

$$H_{t,s},H_{t,u},C_{t,s},C_{t,u}>0$$

The marginal products of the four production inputs are:

$$\frac{\delta G}{\delta h_{u,t}}=(1-\xi)(\xi(\alpha(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}}+(1-\xi)h_{u,t}^{\delta})^{\frac{1}{\delta}-1}h_{u,t}^{\delta-1}$$

$$\frac{\delta G}{\delta h_{s,t}}=(1-\alpha)\xi(\xi(\alpha(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}}+(1-\xi)h_{u,t}^{\delta})^{\frac{1}{\delta}-1}(\alpha(\mu k_t^{\nu}$$
$$+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}-1}h_{s,t}^{\phi-1}$$

$$\frac{\delta G}{\delta k_t}=\xi\alpha\mu(\xi(\alpha(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}}+(1-\xi)h_{u,t}^{\delta})^{\frac{1}{\delta}-1}(\alpha(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}$$
$$+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}-1}(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}-1}k_t^{v-1}$$

$$\frac{\delta G}{\delta e_t}=\xi\alpha(1-\mu)(\xi(\alpha(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}}+(1-\xi)h_{u,t}^{\delta})^{\frac{1}{\delta}-1}(\alpha(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}}$$
$$+(1-\alpha)h_{s,t}^{\phi})^{\frac{\delta}{\phi}-1}(\mu k_t^{\nu}+(1-\mu)e_t^{\nu})^{\frac{\phi}{\nu}-1}e_t^{v-1}$$

The first-order necessary conditions are given by:

$$(1)\ \ \frac{\Psi\theta s}{c_{t,s}} = \lambda_t$$

$$(2)\ \ \frac{(1-\Psi)\theta u}{c_{t,u}} = \lambda_t$$

$$(3)\ \ \frac{\Psi(1-\theta)s}{1-h_{t,s}} = \lambda_t z_t \frac{\delta G}{\delta h_{s,t}}$$

$$(4)\ \ \frac{(1-\Psi)(1-\theta)u}{1-h_{t,s}} = \lambda_t z_t \frac{\delta G}{\delta h_{u,t}}$$

$$(5)\ \ \lambda_t = \beta E_t \lambda_{t+1}\{z_{t+1}\frac{\delta G}{\delta k_{t+1}} + (1-\gamma)\}$$

$$(6)\ \ p_t = z_t \frac{\delta G}{e_t}$$

The log-linearized versions of equations (1)-(6) coupled with the laws of motion for the two shocks and the (log-linearized) aggregate resource constraint yield solutions for the percentage deviations from the steady state for the nine variables $(\lambda, c_u, c_s, h_u, h_s, y, e, k, z, p)'$.

References

[1] 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.

[2] Castañeda, A., Díaz-Giménez, J. and Ríos-Rull, J.V. (2003), "Accounting for Earnings and Wealth Inequality", *Journal of Political Economy*, vol. 111, no. 4, pp. 818-857.

[3] Castro, R. and Coen-Pirani, D. (2005), "Why Have Aggregate Skilled Hours Become So Cyclical since the Mid-1980s?", manuscript, Université de Montreal.

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

[5] Fernández-Villaverde, J. and Rubio-Ramírez, J. (2005), "Estimating Dynamic Equilibrium Economies: Linear vs. Nonlinear Likelihood", *Journal of Applied Econometrics*, 20, 891-910.

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

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

[8] 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.

[9] Kim I. and Loungani P. (1992), "The Role of Energy in Real Business Cycle Models", *Journal of Monetary Economics*, Vol. 29, pp. 173-189.

[10] Klein P. (2000), "Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model", *Journal of Economic Dynamics and Control*, Vol. 24(10), pp. 1405-1423.

[11] 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.

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

[13] Morrison C.J. and Berndt E. (1981), "Short-Run Labor Productivity in a Dynamic Model", *Journal of Econometrics*, Vol. 16, pp. 339-365.

[14] Pindyck R. (1979), *The Structure of World Energy Demand*, The M.I.T. Press, Cambridge, MA.

[15] Polgreen L. and Silos P. (2006), "Capital Skill Complementarity and Inequality: A Sensitivity Analysis", Federal Reserve Bank of Atlanta WP-2005-20.

[16] Prasad E. (1996), "Skill Heterogeneity and the Business Cycle", *Canadian Journal of Economics*, Vol. 29(4), pp. 910-929.

[17] Smith, A. A., (1993), "Estimating Non-Linear Time Series Models Using Simulated Vector-Autoregressions", *Journal of Applied Econometrics*, 8, S63-S84.

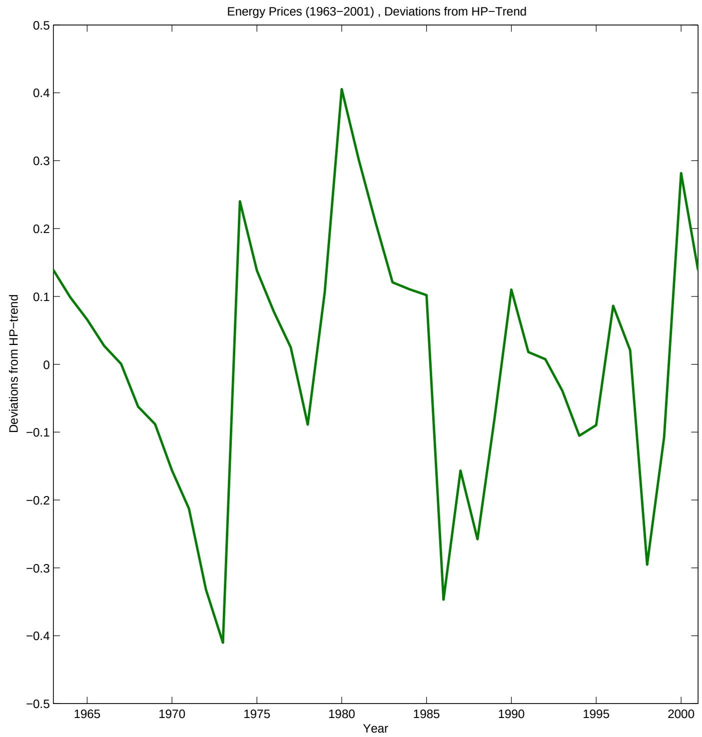

Figure 1: Deviations from an HP-trend of energy prices. US data, annual, 1963-2001.

Figure 2: Deviations from an exponential trend of energy prices and the skill premium. US data, annual, 1963-2001.

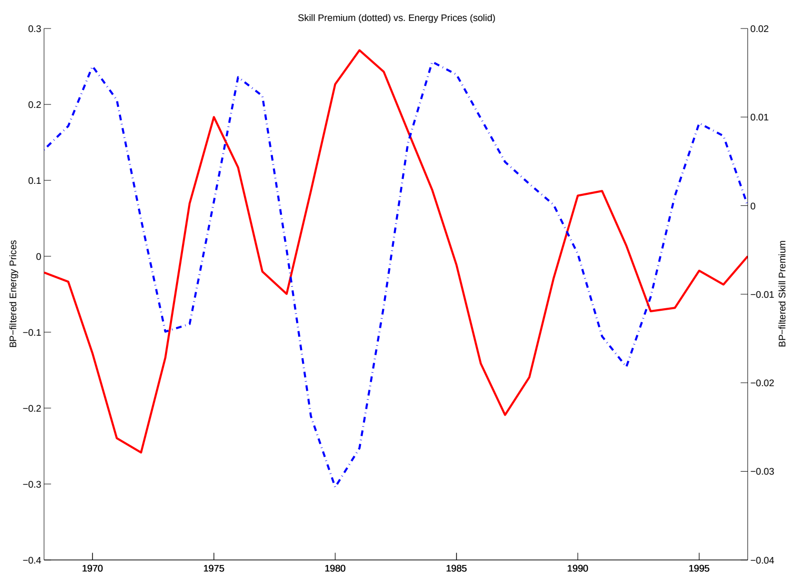

Figure 3: Band-pass filtered energy prices and skill premium. US data, annual, 1963-2001.

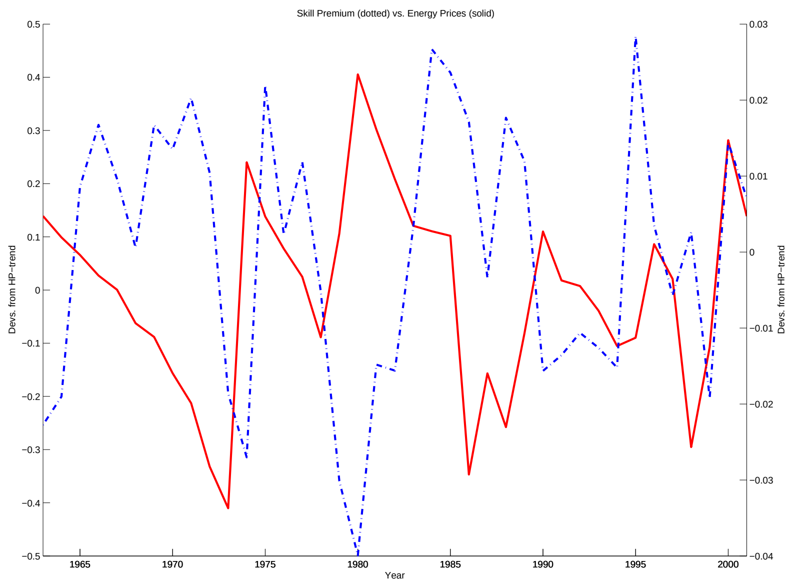

Figure 4: Deviations from an HP-trend of energy prices and the skill premium. US data, annual, 1963-2001.

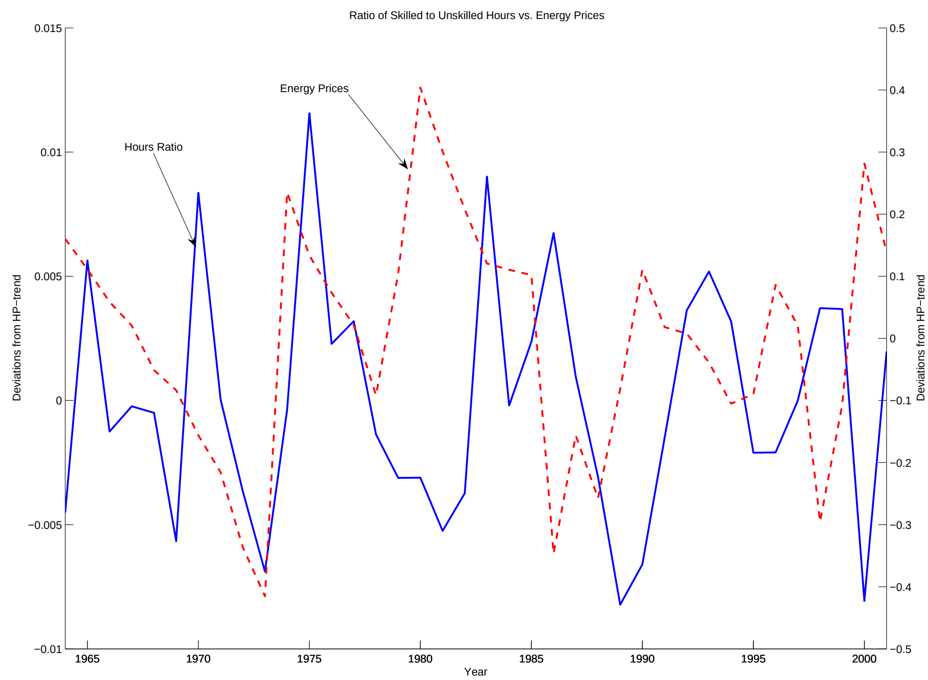

Figure 5: Deviations from an HP-trend of energy prices and the relative hours ratio. US data, annual, 1964-2001.

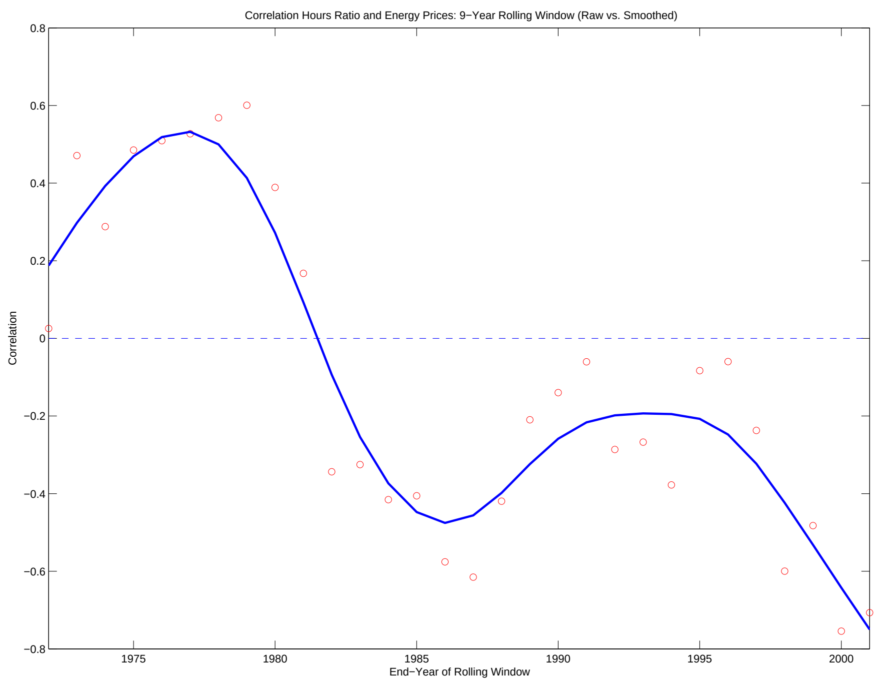

Figure 6: Sequence of 9-year rolling correlations between energy prices and the relative hours ratio. The solid line is the spline-smoothed approximation of the scatter-plot

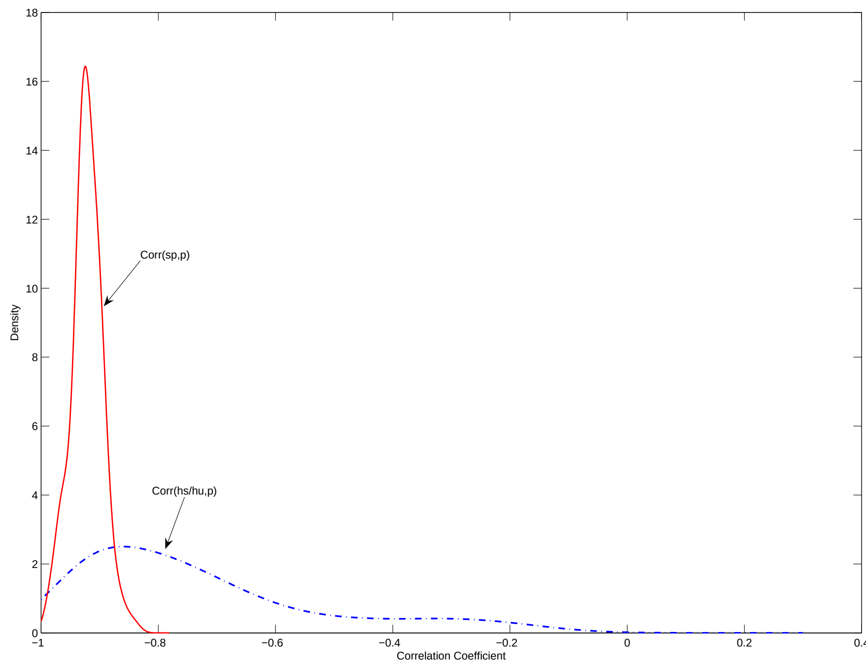

Figure 7: Correlations of the skill premium with oil prices (red solid line) and the hours ratio with oil prices (blue dash-dotted line).

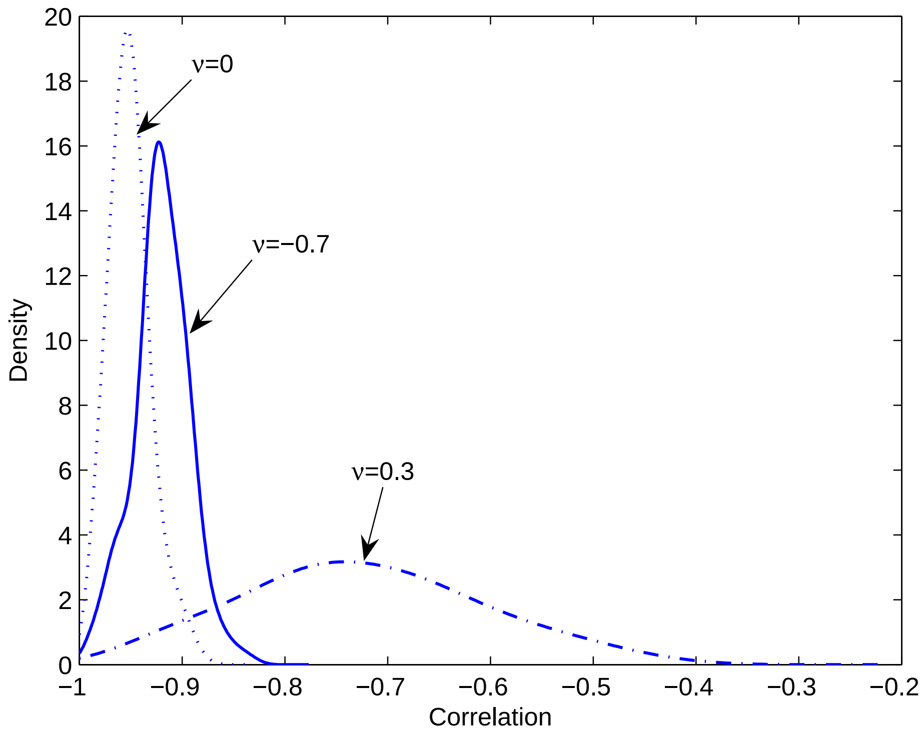

Figure 8: Posterior distribution of the correlation between oil prices and the skill premium in the model for three different values of ν.

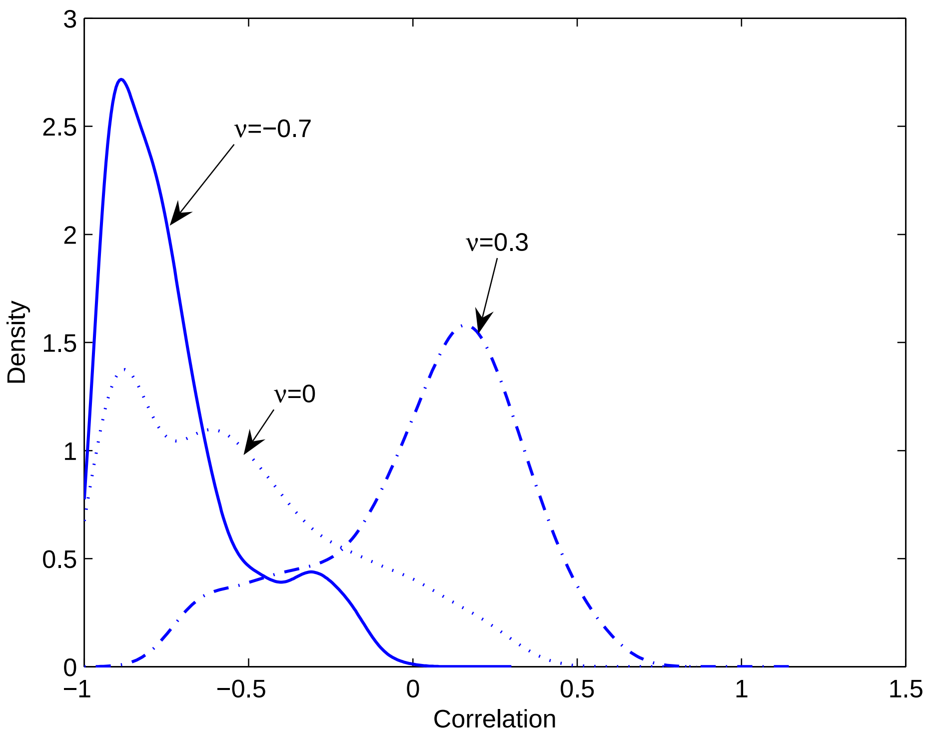

Figure 9: Posterior distribution of the correlation between oil prices and the relative hours ratio in the model for three different values of ν.

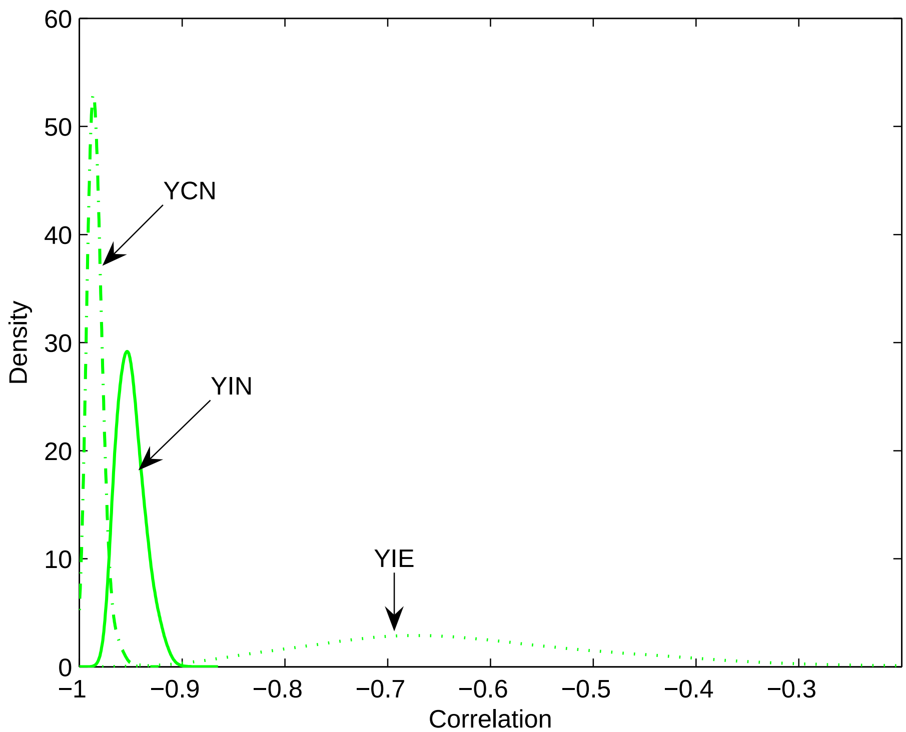

Figure 10: Posterior distribution of the correlation between oil prices and the skill premium in the model for three different sets of estimated parameters.

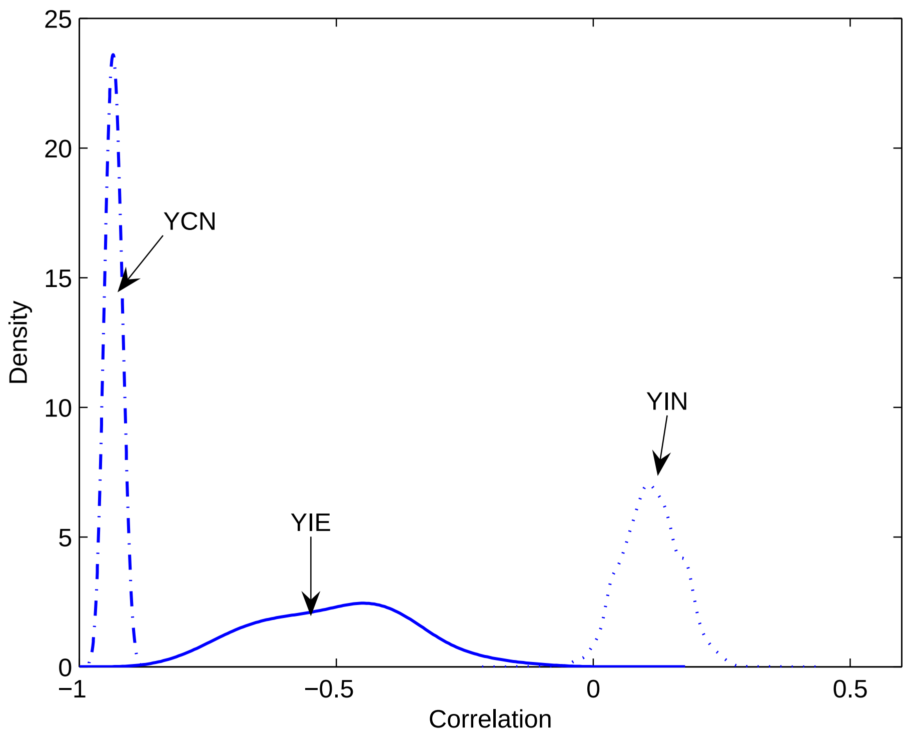

Figure 11: Posterior distribution of the correlation between oil prices and the relative hours ratio in the model for three different sets of estimated parameters.