THE EFFECTS OF ENVIRONMENTAL REGULATION AND ENERGY PRICES ON U.S. ECONOMIC PERFORMANCE

A thesis presented

by

Peter Jensen Wilcoxen

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

December 1988

Chapter 2

2. AN OVERVIEW OF THE MODEL

Because the impact of environmental regulation varied strongly from one industry to another, a disaggregated model must be used to capture its effects. Moreover, the link between regulation and investment means that the model must also incorporate reasonable savings behavior in order to provide useful results. For these reasons, it was most appropriate to develop an econometrically estimated multisector general equilibrium model in which agents have foresight and optimize their behavior over time. The remainder of this section presents the salient features of the simulation model; the complete specification may be found in Appendix A.

Essentially the model consists of two parts, one which determines a single period equilibrium given specific values of certain dynamic variables, and one which links the dynamic variables across time. For convenience, these will often be referred to as the *intra* and *inter*period submodels. The intraperiod part is basically a static general equilibrium system which determines market-clearing prices and quantities for given values of stock variables and expectations. The interperiod segment, on the other hand, determines expectations consistent with the path of stock variables and income flows generated by the intraperiod model. Together, the two parts produce a perfect foresight path of the economy by incorporating into agents' intraperiod decisions expectations about the future that will actually be fulfilled.

In the intraperiod model, production is carried out by 35 industries, each of which produces a primary product and may also make a number of secondary products. The definitions of these sectors are shown in table 2.1. There are 35 corresponding domestic commodities, each representing the primary product of a particular industry. The industries use as inputs commodities and three primary factors: noncompeting imports[1], capital services and labor. Subject to an adjustment discussed below, all industries use a single type of capital and labor. The commodities produced are demanded both by industries for use as intermediate inputs, and by final demand sectors. There are four categories of final demand: consumption, investment, government spending, and net exports. In addition to buying commodities, some of the final demand sectors also

1. Noncompeting imports are goods used in production that are not produced domestically. A typical example is coffee beans.

Table 1.1: The Definitions of Industries

Number	Description
1	Agriculture, forestry and fisheries
2	Metal mining
3	Coal mining
4	Crude petroleum and natural gas
5	Nonmetallic mineral mining
6	Construction
7	Food and kindred products
8	Tobacco manufactures
9	Textile mill products
10	Apparel and other textile products
11	Lumber and wood products
12	Furniture and fixtures
13	Paper and allied products
14	Printing and publishing
15	Chemicals and allied products
16	Petroleum refining
17	Rubber and plastic products
18	Leather and leather products
19	Stone, clay and glass products
20	Primary metals
21	Fabricated metal products
22	Machinery, except electrical
23	Electrical machinery
24	Motor vehicles
25	Other transportation equipment
26	Instruments
27	Miscellaneous manufacturing
28	Transportation and warehousing
29	Communication
30	Electric utilities
31	Gas utilities
32	Trade
33	Finance, insurance and real estate
34	Other services
35	Government enterprises

purchase primary factors. For example, the government purchases labor.

Expectations enter the intraperiod model through the behavior of the household. In particular, households choose their current consumption and labor supply to maximize an intertemporal utility function. This, in turn, requires households to have expectations regarding their future consumption. They do not, however, need to have special expectations about the earnings of the capital stock. Because there are no costs to firms of adjusting their capital stock−other than purchasing the raw capital goods−the value of the firm at any point is just its capital stock multiplied by the current price of new capital goods.

The path of the economy also depends on the rate of growth of a number of stocks. The most evident of these is the capital stock, but there are three others: the stock of government debt held domestically, the government debt held by foreigners, and the stock of foreign debt held domestically. These change from year to year depending on the value of flow variables determined in the intraperiod model.

One unusual feature of the model is that wherever possible, the parameters of behavioral equations were obtained by estimation. This stands in constrast to many general equilibrium models whose parameters are obtained by calibrating the model to a particular benchmark year. Estimation differs from calibration in two respects. First, calibration requires the use of functional forms whose parameters can be determined from a single data point. In contrast, estimation allows any functional form to be used that can be estimated reliably. Second, calibration is not equivalent to estimation (in the usual sense) of the simpler functional forms. For example, the usual estimation procedure for a Cobb-Douglas unit cost function would produce parameter estimates that were equal to the sample mean cost shares of the various inputs. Calibration, however, produces the cost shares for a particular year. To the extent that the benchmark year is unusual, calibration will produce results substantially different from estimation of the same functional forms. This is the root of the problem addressed by Higgs (1986).

For most of the behavioral equations, transcendental logarithmic functions were used because of their ability to capture a wide variety of interactions between inputs, avoiding the unitary elasticity of substitution implicit in Cobb-Douglas functions, and allowing more flexibility than the Constant Elasticity of Substitution function. These benefits are obtained at some cost, however, as translog functions usually have a large number of parameters, and it may not be possible to obtain enough data to be able to estimate them all.

Details of the estimation of all the behavioral equations appears in Appendix B. Some of the parameters were estimated specifically for this model, while others were taken from the earlier work of various people on related models. The parameters alone, however, are not enough to specify the model completely because not all of the model's variables are described by behavioral equations. The most obvious example is tax rates, which are treated as exogenous. Many other variables of this type appear in the model, and their values had to be obtained before the model could be solved. The sources and methods used to construct the exogenous data are also described in Appendix B.

Finally, one additional feature that distinguishes this model from others is the treatment of time in some of the behavioral equations. In particular, time itself does not enter the production function, but rather a logistic function of it does. This specification was used because functions linear in time are inappropriate for use in this model because it will be simulated far beyond the end of the sample period used in estimation. As a practical matter, if the production has biased, or even unbiased technical change that is linear in time, the model will exhibit bizarre behavior far in the future. The only exception to this occurs in the unlikely event that technical change is unbiased and the same for all industries. If the former is not the case, technical change will eventually drive more and more input cost shares negative; if the latter does not hold, over time the economy will become dominated by the fastest growing sectors. The most important problem, however, is conceptual. Given only information from the sample period, it is hard to justify assuming that technical change will continue to be like it is now a long time in the future.

Replacing time with a logistic function eliminates these problems because technical change eventually disappears as time gets far into the future. As a result, cost shares become stable, and unbiased technical change goes to zero in all industries. This amounts to assuming that the technical change observed in the sample period resulted from a gradual, but essentially one-time, change from past production functions to those of the future. At a fundamental level, this is similar to the linear-time assumption above: both are nothing more than speculation about the path of technology in the future. As such, neither formulation is more likely to be "true" than the other, in the sense that it is likely to be observed. The logistic formulation produces reasonable behavior near the sample period while also providing a steady state far in the future. Moreover, it has the advantage of being agnostic about the direction of future technical progress. The actual production function used is discussed in detail in a subsequent section.

Time also enters the household model in a novel way. Specifically, some of the parameters that determine how the household decides on the mix of consumption and leisure are allowed to be functions of time. This was used to account for the rapid entry of women into the labor force over the 1970's and 1980's. Without this feature, the model would consistently overstate labor supply in the past and understate it in the future. Having outlined the important and unusual features of the model, it's now useful to discuss some of them in more detail.

2.1. The Intertemporal Model

The foundation of any dynamic model is its treatment of savings and investment. Early efforts, such as Goettle and Hudson (1981), derived savings from current income and consumption, and then used that to determine investment. In contrast, internal adjustment cost models, such as Wilcoxen (1987), provide a rigorous, forward-looking treatment of investment while assuming passive savings behavior.[2] Other models, such as Goulder and Summers (1987), combine internal adjustment costs with savings behavior determined by an intertemporal consumption decision. Finally, the model presented here is representative of a fourth category−it has intertemporal consumption behavior and a rigorous treatment of investment, but no adjustment costs. To see how it works, it will be convenient to discuss the investment and savings aspects in turn.

2. Wilcoxen (1987) presents an investment model suitable for a small open economy that can borrow at a fixed world interest rate.

The basis of the investment model is relationship between the cost of new capital goods and the present value of the stream of returns they will earn in the future. As long as capital goods cost less than the returns they'll earn, the net benefit of additional investment is positive, so investment will increase. On the other hand, if returns fall below the price of new assets, investment will diminish. In equilibrium, then, the price of capital goods must be exactly equal to the returns earned on them. The value of installed capital goods can be determined by the arbitrage equation between returns on government bonds and equity:

$$rV = D + \frac{dV}{dt} \tag{1.1}$$

where r is the rate of return on bonds, V is the value of capital, D is the dividend earned on capital, and dV/dt represents capital gains. Since there is a single capital good in the model, the dividend earned on it is simply the sum of rental payments from all sectors.

At this point, the absence of internal adjustment costs comes into play. Without adjustment costs, the price of a new capital good is always equal to the cost of producing it. Thus, the following expression must hold:

$$V = P^K K \tag{1.2}$$

where K is the capital stock, and P^K is purchase price of investment goods. This means that investment is very elastic: when expected returns on capital rise, the desired amount of investment will go up until either the returns are lowered, or the price of capital goods increases from general equilibrium effects. With elastic investment, the actual amount of it done will depend heavily on the supply of savings.

Savings are determined by households as a result of their consumption and labor supply decisions. Although certain adjustments are made to account for aggregation, the model has essentially one infinitely-lived consumer who chooses a path of consumption to maximize an intertemporal utility function. Moreover, the consumer has perfect foresight and can correctly predict future prices and incomes. The objective function used is shown below:

$$U = \sum_{t=0}^{\infty} N_0 \prod_{s=1}^{t} \left(\frac{1 + n_s}{1 + \rho} \right) \ln F_t \tag{1.3}$$

where F_t is a per capita aggregate of goods and lesiure consumed in period t (called "full" consumption below), ρ is the rate of time preference, N_0 is the initial population, and n_s is the population's growth rate. Since households consider all future earnings when making current consumption decisions, this model will exhibit permanent income behavior. In particular, temporary changes in income will be reflected strongly in savings and weakly in consumption.

Maximizing utility subject to the intertemporal budget constraint produces the Euler equation shown below. It relates consumption in two arbitrary adjacent periods, 1 and 2, which can be thought of as "today" and "tomorrow".

$$P_1^F F_1 = P_2^F F_2 \left(\frac{1 + \rho}{1 + r_2} \right) \tag{1.4}$$

P_1^F and P_2^F are the prices of full consumption in the two periods. This expression determines the slope of the optimal consumption path, while the absolute level is set by the budget constraint. The interest rate plays a crucial role: if r is larger than ρ, full consumption will be rising over time; if it's smaller, consumption will fall. Moreover, increases in r make the path steeper, lowering consumption today relative to tomorrow. This is the basis for the supply of savings in the model: at constant current income, an increase in r will lower consumption and increase savings.

The actual amount of investment is determined by the interaction of the savings and investment functions with the path of interest rates. In this respect the model is similar to that of Abel and Blanchard (1983). For an expected path of interest rates, the savings mechanism described above determines the amount of investment. If the results differ from what was implied by the investment function, actual interest rates will be higher or lower than expected. This, in turn, will cause the savings supply to change. In equilibrium, the path of investment must satisfy both the savings and investment functions simultaneously.

To complete the discussion of investment, it's important to describe the capital stock itself. An unusual feature of the model is that both consumer durables and owner occupied housing are treated as capital. This means, for example, that household demand for automobiles is part of investment, not consumption, and contributes to the growth of the capital stock. Moreover, household consumption of housing services is treated as a demand for capital, not for the output of the real estate sector (as is done by the BEA in its input-output accounting). This approach correctly implements the theoretical symmetry between household and other types of capital.

Finally, new capital goods are formed out of the commodities making up the investment final demand column according to an estimated production function that allows substitution between inputs. This stands in contrast to the fixed-coefficients approach used in many models in which the commodity composition of investment is set according to a vector of base year shares. The functional form used is described in detail in Appendix A, and the estimated parameters are discussed in Appendix B.

2.2. Production and Technical Change

Production was modelled using nested transcendental logarithmic unit cost functions to allow the maximum feasible flexibility in the interactions between inputs. Before discussing the actual specification, however, it's useful to note a few of the important features of translog cost functions in general. Consider an industry purchasing a vector of n inputs whose prices are given by vector P. A typical translog unit cost function would be the following:

$$\ln C = \alpha_0 + \alpha_P{}' \ln P + \frac{1}{2} \ln P' \beta_{PP} \ln P \tag{1.5}$$

where α_0 is a scalar parameter, α_P is a vector of parameters of length n, and β_{PP} is an $n \times n$ array of parameters. By Shepard's Lemma, the demand for factor i per unit of output can be derived from the cost function by differentiation:

$$\frac{X_i}{Q} = \frac{\partial C}{\partial P_i} \tag{1.6}$$

Multiplying both sides by $\frac{P_i}{C}$ produces an expression giving the share of input i, ω_i, in total cost:

$$\omega_i = \frac{X_i P_i}{QC} = \frac{\partial C}{\partial P_i} \frac{P_i}{QC} = \frac{\partial \ln C}{\partial \ln P_i} \tag{1.7}$$

Applying this to the translog production function produces a vector ω of input shares which have the following formula:

$$\omega = \alpha_P + \beta_{PP} \ln P \tag{1.8}$$

This relationship is vital to the estimation of translog production functions because it is observable; it is also of considerable use in the simulation model because it can be used to compute input-output coefficients. Unfortunately, the number of parameters to be estimated is of the order of $1 + n + n^2$, so it is usually impossible to estimate a function with more than a handful of inputs.[3] Because of this, it is often necessary to make a number of separability assumptions to

3. It is actually somewhat smaller than this figure because of constraints on the parameters.

allow the overall function to be estimated as a collection of nested aggregates. The problem is particularly apparent in estimation of the industry production functions which have 38 inputs (35 intermediate goods and 3 primary factors). In practice it is difficult to estimate functions with more than about 5 inputs, so an extensive set of separability assumptions and nesting was required.

The actual cost function used is described in detail in Appendix A, but it's useful to discuss it briefly here to highlight some of its unusual features. A nested tier structure was employed to group the 38 inputs into aggregates. At the top level, output was produced using capital, labor, energy and materials (KLEM). Energy and materials were, in turn, nested functions of the 35 intermediate goods. Technical change was allowed to enter at the KLEM level, but not at lower tiers. The form of the top tier cost function was the following, where P is a vector of prices of capital, labor, energy and materials:

$$\ln C = \alpha^0 + \alpha^{P\prime} \ln P + \frac{1}{2} \ln P' \beta^{PP} \ln P + \alpha^T g(t) + \ln P' \beta^{PT} g(t) + \frac{1}{2} \beta^{TT} g^2(t) \qquad (1.9)$$

Differentiating 2.9 with respect to the log of input prices gives the vector of cost shares shown below:

$$\omega = \alpha^P + \beta^{PP} \ln P + \beta^{PT} g(t) \qquad (1.10)$$

When the functions were estimated, all of the parameter restrictions implied by theory were imposed. These are discussed in detail in Appendix B, and include the integrability conditions: homogeneity, product exhaustion, symmetry, nonnegativity and local concavity. For a discussion of these properties, refer to Jorgenson (1986). By imposing these restrictions, the resulting cost functions are fully consistent with the theory of production.

This formulation allows for both neutral and biased technical change. The parameters α^T and β^{TT} contribute to neutral change by lowering the cost of output without changing the shares of the inputs. In contrast, the β^{PT} term allows for biased technical change since the cost shares change over time, even if prices are constant.

These biases present a potential problem. When the model is simulated far into the future, the value of t will become very large. Without choosing the function $g(t)$ carefully, it is possible that certain cost shares will be driven negative over time. Other models, such as Goettle and Hudson (1981), used $g(t) = t$, and some cost shares become negative very rapidly. To avoid this problem, a logistic formulation was used for $g(t)$ as shown below:

$$g(t) = \frac{1}{1 + e^{-\mu(t-\tau)}} \qquad (1.11)$$

where μ and τ are parameters estimated separately for each industry (see Appendix B). The key feature of this is that $g(t)$ becomes constant at $t \to \infty$. This is of vital importance because the technical change biases are attenuated as t becomes large. As a result, the cost shares become constant, so the problem above is eliminated.

The logistic specification also has consequences for the neutral component of technical progress. Differentiating the log of the cost function with respect to time gives the rate of technical change:

$$v = (\alpha^T + \ln P' \beta^{PT} + \beta^{TT} g(t))\dot{g}(t) \qquad (1.12)$$

As t becomes large, $\dot{g}(t)$ goes to zero, so neutral technical change disappears. This is important because the initial rates of progress differ widely across industries. Without an attenuation of technical change, some sectors' costs would drop so much that they would come to dominate the economy.[4] Using the logistic eliminates this problem by bringing all sectors to constant, albiet zero, rates of progress in the future.

The gradual elimination of technical change is one of two features that together cause the model to have a steady state. The second is the forecast rate of growth of the labor force. As discussed in Appendix B, both the rate of population growth and the rate of increase in educational attainment are forecast to decline to zero over the next 100 years. This means that the rate of growth of the effective labor force will eventually decline to zero. Without technical change or a growing labor force, the economy will eventually attain a stationary state in which there is no income growth at all. This provides the model with a steady state.

Overall, using a logistic function for $g(t)$ allows the model to capture effects in the sample period well, while also producing reasonable behavior far in the future. It may not be completely satisfactory in all circumstances, but it is a distinct improvement over other common resolutions to the problems above. For example, one such approach is to have all growth originate from a single source, usually the labor force. This provides the model with a balanced growth equilibrium, but at the cost of eliminating the large differences in growth between industries[5].

Finally, the aggregate rate of productivity growth can be computed from the industry rates computed by the model using the approach developed by Jorgenson, Gollop and Fraumeni (1987). Suppose the economy can be represented by an aggregate production function that depends only on capital, labor and time.[6] For this assumption to be valid, it can be shown that each sector must have a value added function which differs from the aggregate by at most a multiplicative constant. Moreover, since all technical change must come through value added, the industry rates of productivity growth have to be transformed in the following way.

Consider an arbitrary unit cost function that depends on the prices of capital, labor, intermediate goods and time:

4. Moreover, the model would never reach either a steady state or a balanced growth equilibrium. Since the usual method of solving perfect foresight models is to impose a transversality condition derived from one of these two outcomes, their absence would pose a formidable difficulty.

5. As shown by Jorgenson (1988) for the period 1948-1979, average industry growth rates range from less than zero to more than six percent.

6. At the aggregate level intermediate goods all cancel out.

$$C = C(P_K, P_L, P_X, t) \tag{1.13}$$

where P_X is the price of intermediate goods. Growth rates may be obtained by logarithmic differentiation:

$$\partial \ln C = S_K \partial \ln P_K + S_L \partial \ln P_L + S_X \partial \ln P_X + v^A \tag{1.14}$$

The coefficients, such as S_K, are the value shares of each input in total costs.[7] The rate of technical change obtained in this way, v^A, corresponds to that found in the model. In contrast, to meet the requirements of the aggregate production function, all productivity growth must come through value added, and be separable from intermediate input. This requires the individual industry cost functions to have the following form:

$$C(P_V(P_K, P_L, t), P_X) \tag{1.15}$$

Logarithmic differentiation of this produces the expression shown below:

$$\partial \ln C = S_V(S*_K \partial \ln P_K + S*_L \partial \ln P_L + v^B) + S_X \partial \ln P_X \tag{1.16}$$

In this expression, the coefficients are again cost shares; those marked with an asterisk are shares in the value added function. This can be simplified by noting that $S_V S*_K$ is just S_K, and that a similar fact is true for labor. Thus, the equation can be rewritten as shown:

$$\partial \ln C = S_K \partial \ln P_K + S_L \partial \ln P_L + S_V v^B + S_K \partial \ln P_X \tag{1.17}$$

7. This can be demonstrated by applying Shepard's Lemma.

Comparing this with the equation derived above for the more general case shows that the two rates of technical change are related as shown:

$$v^B = \frac{v^A}{S_V} \tag{1.18}$$

This makes perfect sense−to obtain a given rate of output growth, technical change must be larger if it is confined to augmenting value added rather than all inputs.

After the industry level rates of productivity growth are transformed in this way, aggregate productivity growth can be found by computing a weighted sum of them, where the weights are each sector's share in total value added.[8] Calculating aggregate productivity growth will make it possible to assess the effect of different policies on the rate of technical change.

2.3. The Special Case of Oil Extraction

As the role of oil prices was a central object of investigation and the crude oil industry has a number of unusual features, it was modelled somewhat differently from the other sectors. In particular, the supply elasticity of petroleum is quite low, especially in the short run, so it is inappropriate to represent it by a constant returns to scale cost function. One reason the elasticity is so low is that the industry's primary capital stock, underground reserves, can be changed only very slowly as it takes years to find and develop new fields (Christiansen and Reister, 1988). This suggests that a resonable approach to modelling oil production is to fix the industry's capital stock, at least in the short run. Doing so would reduce the elasticity of supply considerably, and for a plausible reason.

There are two slightly different methods available for fixing the oil industry's capital stock. One approach is to modify the sector's cost function to reflect the capital constraint. This would

8. This assumes there are no changes in productivity from transferring factors from one sector to another. Jorgenson, Gollop and Fraumeni suggest, however, that this effect can be important during certain periods.

produce the necessary upward sloping supply curve, but it would also eliminate the possibility of using the nonsubstitution theorem to determine the industry’s price. Moreover, it would destroy any degree of symmetry in the treatment of different sectors. A second tactic is to maintain the original cost function, but limit the amount of capital available to the sector. In this situation, firms in the industry would take the rental price of capital as given when deciding how much of it to use. The rental price would then be determined by the industry’s demand for capital relative to a fixed supply of it. Using this method, the nonsubstitution theorem could still be applied to determine the price of output. It also preserves a high degree of symmetry in the treatment of industries. For these two reasons, the capital constraint was imposed using the second approach.

With the capital stock fixed, the sector’s effective supply elasticity can be computed from its unit cost function as follows. The share of costs attributed to capital, ω_k, is given by the expression

$$\omega_k = \frac{P_k K}{P_o Q} \tag{1.19}$$

where P_k is the rental cost of capital, P_o is the price of output, and Q is the quantity of output. Logarithmic differentiation of 2.13 produces the following:

$$\partial \ln \omega_k = \partial \ln P_k + \partial \ln K - \partial \ln P_o - \partial \ln Q \tag{1.20}$$

Differentiating the log of the cost function holding the prices of all inputs constant except for the capital stock yields:

$$\partial \ln P_o = \omega_k \, \partial \ln P_k \tag{1.21}$$

Finally, differentiating 2.10 gives

$$\partial \ln \omega_k = \frac{1}{\omega_k} \beta_{kk} \partial \ln P_k \tag{1.22}$$

Substituting 2.15 and 2.16 into 2.14, and holding capital constant produces the elasticity of output with respect to its price:

$$\eta = \frac{\partial \ln Q}{\partial \ln P_o} = -\left(\frac{\beta_{kk} - \omega_k + \omega_k^2}{\omega_k^2} \right) \tag{1.23}$$

For crude oil production, β_{kk} is about .08, while ω_k is 46%. Inserting these values into expression 2.17 and working out the calculation shows that η is around 0.8, a value that is roughly consistent with other findings, such as those reported by Kennedy (1974). This specification should perform well for experiments in which the price of oil does not change a lot, or only changes temporarily. For large, permanent changes it may be inappropriate, as it is likely that drilling would respond in the long run. Further analysis of this issue is presented in Chapter 4.

One other feature of the oil sector is of particular interest: the elasticity of substitution between imported and domestic petroleum. It is useful to derive an expression for this value in terms of the parameters of the model. By definition, the shares of foreign oil in total supply is

$$\omega_f = \frac{X_f P_f}{P_o Q} \tag{1.24}$$

The domestic share is similar. Solving for the two inputs, taking the ratio and applying logarithmic differentiation gives the following

$$\partial \ln(X_f/X_d) = \partial \ln \omega_f - \partial \ln \omega_d - \partial \ln(P_f/P_d) \tag{1.25}$$

From Appendix A, the value share of foreign oil in total supply has the form

$$\omega_f = \alpha_f + \beta_{fd} \ln P_d + \beta_{ff} \ln P_f \tag{1.26}$$

Since there are only two inputs it must be true that $\beta_{ff} = -\beta_{fd}$, so this can be rewritten as

$$\omega_f = \alpha_f + \beta_{fd} \ln(P_d/P_f) \tag{1.27}$$

Taking logs and differentiating produces

$$\partial \ln \omega_f = \frac{1}{\omega_f} \beta_{fd} \, \partial \ln(P_d/P_f) \tag{1.28}$$

Finally, inserting this and the corresponding expression for ω_d into 2.19 and rearranging produces the elasticity of substitution of foreign and domestic crude oil:

$$\sigma_{fd} = -\left(\frac{\beta_{fd} + \omega_f \omega_d}{\omega_f \omega_d} \right) \tag{1.29}$$

When β_{fd} is zero, production of the composite is Cobb-Douglas and the corresponding elasticity is one. As β_{fd} becomes positive, foreign and domestic oil will be better substitutes, while they will be worse if it is negative.

As discussed in Appendix B, for most sectors this elasticity was estimated from data. For crude oil, however, institutional characteristics of the market made this impossible. In particular, import quantity controls were in place for most the sample period (Greenberger, 1983), so buyers were not able to substitute freely between goods of different origin. The domestic price was kept artifically low, and quotas were used to allocate it to refiners. The seriousness of this problem is evident from the results obtained if estimation is attempted: the elasticity of substitution is on the order of .5, exceedingly small for a commodity as apparently homogeneous as crude oil. Unfortunately, it is not sufficient simply to set the elasticity to a large number, because that would not capture the actual behavior of the industry during the sample period. The domestic price of crude would follow the foreign price much more closely than it actually did.

There is no completely satisfactory resolution to this problem. Making the elasticity large will distort the behavior of the domestic industry during the sample period. On the other hand, making it very small is hard to justify given the chemical similarity of the products. As a compromise, the elasticity was set to one. This has the effect of making the aggregation of foreign and domestic oil Cobb-Douglas, so the shares of the sources in total supply will be constant over time. The role of this parameter is explored further in Chapter 4, where some sensitivity results are presented.

2.4. The Government and Current Account Deficits

Almost all the variables in the model were determined by behavioral equations, equilibrium conditions, or accounting identities. Regrettably, there were two exceptions: the government and current account deficits. It was beyond the scope of this study to develop structural models of these, so they had to be treated differently. The government's budget was handled as follows: tax rates and the deficit were set exogenously, while spending was allowed to adjust to satisfy the budget constraint. The tax rates were taken from actual data during the period 1974-1985, and thereafter held at their 1985 values. For the deficit, actual data was used through 1985; after that it was forecast to close gradually over the next twenty years. Details of the projections are contained in Appendix B.

This not the only partition that could have been used. It would also have been possible to make government spending exogenous and taxes endogenous using a lump sum tax. Unfortunately, this would require projections of expenditure, in addition to the deficit, which would be very difficult to forecast with any degree of confidence. In contrast, fixing tax rates at their 1985 values will undoubtedly turn out to be false, but is unlikely to be as far from the truth as any expenditure forecasts. The final possible partition would be to set expenditure and tax rates exogenously, and make the deficit endogenous. Given recent experience, this formulation seems most like the truth. However, an endogenous nonzero deficit would prevent the model from ever attaining a steady state, so this partition could not be used.

The current account deficit presented a similar problem. In this case, either the exchange rate had to be made exogenous and the balance of trade endogenous, or the reverse. As with the government deficit, however, allowing the current account to be endogenous is inconsistent with attainment of a steady state. For this reason, the exchange rate was made endogenous and the current account was set exogenously. Before 1985 it was taken from actual data, while later years were forecast. In the projections, the deficit was forced to zero by the year 2000. After that, the current account was moved to surplus for a number of years to bring net domestic holdings of foreign assets back to a small positive value. With an endogenous exchange rate, as the the current account climbed back to surplus, a temporary devaluation also occurred.

This choice of partitioning had a marked effect on simulations in which the current account would ordinarily be expected to change. In particular, the exchange rate was forced to change as much as necessary to produce the specified trade balance under the new conditions. This has a clear interpretation: the new exchange rate shows what would have to hav happened to the terms of trade to produce the given deficit under the new conditions. If, for example, the foreign price of oil were suddenly increased, the trade deficit would not be allowed to grow, so the exchange rate would have to depreciate enough to keep it from rising.