NBER WORKING PAPER SERIES

ENERGY TAXES AND AGGREGATE ECONOMIC ACTIVITY

Julio J. Rotemberg
Michael Woodford

Working Paper No. 4576

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December, 1993

We would like to thank Jim Costain for his untiring research assistance, Jim Poterba for his comments and the NSF for research support. This paper is part of NBER's research programs in Economic Fluctuations and Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4576
December 1993

ENERGY TAXES AND AGGREGATE ECONOMIC ACTIVITY

ABSTRACT

This paper shows that the output losses from energy taxes are significantly larger than usually computed when due account is taken of imperfect competition among energy using firms. Even with perfect competition among these firms, the loss in GNP is of the same order of magnitude as the revenue raised by these taxes. However, in the presence of imperfect competition the output losses are much higher. There are particularly large transitory losses in the immediate aftermath of energy price increases when firms act as implicitly colluding oligopolists. These losses become considerably smaller if energy taxes are phased-in. We also show that taxes that affect only household consumption of energy have much smaller effects. In particular, for the empirically plausible parameter values we consider, such taxes have no effect on employment or output in the non-energy sector.

Julio J. Rotemberg
Department of Economics, E52-432
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02139
and NBER

Michael Woodford
Department of Economics
University of Chicago
1126 E. 59th Street
Chicago, IL 60637
and NBER

As part of his address to a joint session of Congress on February 17, 1993, President Clinton proposed a broad-based energy tax, as a central part of his plan to reduce the size of the U.S. government budget deficit. Had this "BTU tax" been enacted, crude oil would eventually have been subject to approximately a 21% tax, coal to a 25.7% tax and natural gas to a 16% tax. Somewhat lower taxes would have applied to hydroelectricity and nuclear power. The political resistance to this energy tax was, however, intense, and when the dust settled, all that was enacted was about a 4% tax on gasoline.

One of the reasons advanced for resistance to the energy tax was concern about its impact on production and employment in U.S. industry. Indeed, existing studies of the effect of carbon taxes (Goulder (1992, 1993a, 1993b), Jorgenson and Wilcoxen (1993)) suggest that the reductions in GDP caused by these taxes are comparable to the amount of revenue they raise. To demonstrate that this is more onerous than the losses caused by other taxes, these authors show that GDP still falls substantially even if the revenues from the carbon tax are used to reduce existing labor income taxes. [1] That energy taxes are so deleterious may seem surprising, since energy consumption is a relatively small fraction of GDP. But the share of energy costs in total costs does not affect the analysis because the small share of spending on energy also reduces proportionally the revenue raised by a given *ad valorem* tax rate. Neither does our analysis hinge on the fact that, in practice, other inputs are used to produce energy. Thus the cost of energy taxes we discuss is unrelated to Diamond and Mirrlees' (1971) proof that it is inefficient to tax intermediate inputs; energy is actually a raw material in our model since we neglect extraction costs. Rather the cost of energy taxes results from the fact that, unlike other raw materials such as labor, energy is relatively elastically supplied. As a result, the quantity of the energy input falls substantially in response to a tax, instead of the factor price simply being forced down.

In this paper, we argue that the contractionary effects of energy taxes on energy-using industries are even larger than is usually computed, once due account is taken of imperfect competition in those industries. [2] The presence of imperfect competition implies that the price of output is above the marginal cost of production. Thus the social benefit from increasing output by one unit exceeds the social cost of

[1] Of course, an appropriately structured energy tax also has a benefit that other kinds of taxes do not, which is the provision of a disincentive for activities with harmful external effects. This, rather than the search for additional sources of government revenue, is the main reason for recent discussion of "carbon taxes". From this point of view, an energy tax can actually improve efficiency. Because we do not here attempt an overall evaluation of the welfare consequences of an energy tax, we do not attempt to quantify such effects. For an attempt to do so, see Goulder (1993b)

[2] Judd (1993) shows that imperfect competition also affects the optimal tax on capital income. His analysis differs from ours because capital goods are intermediate inputs whereas we treat energy as a raw material.

doing so. This wedge implies that a reduction in output has more deleterious welfare consequences in the presence of imperfect competition. Thus, the preexisting distortion due to the lack of perfect competition raises the welfare costs of any particular output reduction, whatever its origin. Welfare costs are not out main focus here, however. We study instead the degree to which output falls and show that this too is larger with imperfect competition.

The reason is twofold. First of all, imperfect competition implies that the marginal product of any factor, including energy, exceeds its price. This means that the reductions in energy and other inputs that result from energy taxes reduce GNP by more than one would estimate simply based upon these inputs' measured cost shares.

Secondly, if the tax change increases the degree of market power of firms in their product markets, they increase the extent to which they mark up their prices relative to their marginal costs, which results in a contraction of the equilibrium level of production, just as if a tax on inputs had increased those marginal costs. We show that even a very small increase in market power can have a large effect upon the predicted output decline, because the markup increase is like a tax on all inputs, and not just energy. We also show that a particular model of endogenous markup determination (the model of oligopolistic pricing previously used in Rotemberg and Woodford (1991, 1992, 1993)) can imply a temporary increase in market power following an energy tax increase, though the effect is transitory even in the case of a permanent tax increase. We also show that this effect is even stronger if one allows for uncertainty about the permanence of the tax change.

We also show that allowing for imperfect competition has important consequences for evaluation of the relative merits of alternatively structured energy taxes. In particular, we show that gradual phase-in of an energy tax mitigates the contractionary effects in the short run, to an even greater extent than revenues are reduced over that same period; and this effect is even more pronounced when imperfect competition is taken into account.

Our method is to numerically solve a calibrated general equilibrium simulation model, under alternative assumptions about product market structure. Our model decomposes energy into energy purchased directly by households and energy bought indirectly via the purchase of other produced goods. This allows us to analyze the difference between taxes on all energy use and taxes on directly consumed energy.

This paper is related to Rotemberg and Woodford (1993), where we considered the ability of a similar range of alternative models to explain the large declines in U.S. output that followed pre-1980 increases in the price of oil. We showed that it was easiest to explain these contractions of output, as well as the simultaneous declines in real wages, if one viewed firms as not only having market power but as implicitly collusive. [3] The numerical calibration of the "variable markup" model considered here matches that of the model shown in the previous paper to best fit the observed effects of oil price shocks. This gives us some reason to suppose that imperfectly competitive effects of the size assumed in our simulations may actually be present in the U.S. economy.

Section 1 sets the stage by describing the U.S. energy market. Section 2 discusses the behavior of the firms that use this energy to produce final output. In this section, we also give an intuitive explanation for the importance of imperfect competition in determining the output losses caused by energy taxes. Section 3 then describes the rest of our simulation model. Sections 4 and 5 then present the model's numerical predictions regarding, respectively, the long-run and short-run effects of an unexpected permanent increase in energy taxes. In section 6 we take up the effect of predicted changes in energy taxes. We thus consider both the effect of phased-in taxes as well as the effects of taxes that are expected to be repealed. Finally, section 7 concludes.

1 The U.S. Energy Market

Four types of products account for the vast bulk of energy consumption. These include coal, natural gas, petroleum products and electricity. For our purposes, we wish to obtain an energy aggregate. One common approach is to add together the BTU's contained in all four sources of energy. This would make sense if the products were perfect substitutes in the sense that a BTU from one source is as useful as a BTU from another. However, in practice, the price per BTU is rather different for different sources of energy. In particular, it is higher for oil than for coal. For that reason, our aggregate is obtained by adding together the expenditure on these four products. This too is strictly appropriate only if the products are perfect substitutes. However, it allows the BTU's from one source to be less useful than those from another.

[3] The same model of oligopolistic pricing is shown in Rotemberg and Woodford (1991) to be useful in explaining cyclical variations in real wages, and in Rotemberg and Woodford (1992) to be useful in explaining the effects of military purchases on real wages.

The aggregation of these four energy sources is complicated by the fact that coal, gas and petroleum products are used in the generation of electricity (though some electricity is also generated from other sources). It would thus be incorrect to simply add together the values of coal, natural gas, petroleum products and electricity sales. What we do instead is to count only the coal, natural gas and petroleum products that are not sold to electric utilities.

Table 1 presents data on the sales of these four products in 1989. Most of coal is used for electric generation. We valued the 100 million metric tons that are consumed in other sectors at the average CIF price paid by electric utilities, namely $30.43 per ton. To value both domestically produced and imported crude oil, we used the average import price of $16.54 per barrel. Electric utilities do not use crude oil directly. Rather, they buy a combination of different petroleum products. Over half of these are made up of residual fuels whose average price was $18.65 per barrel. We assigned this price to the entire volume of petroleum products purchased by electric utilities. Being higher than the price of crude oil, electric utilities are effectively also purchasing some of the value added of the refining sector. This does not pose any conceptual difficulties since we add the entire value added of the petroleum and coal products sector to our aggregate.

In the case of natural gas, we start with the revenues of the industry. [4] We then subtract the gas purchased by electric utilities using the average price paid by them for natural gas. [5] Finally, we add the total revenues by electric utilities to our aggregate. We conclude that energy consumption in 1989 was equal to 365.4 Billion dollars, or about 6.6% of GDP. Of this, imported oil accounts for $62.3 billion, or .17 of the total, and 1.1% of GDP.

We have less accurate data for the breakdown of energy use between direct household use and non-energy production. In the case of electricity revenue, we know that approximately 1/3 comes from residential sales. In the case of the gas sector, we know that residential sales account for $25.4 billion in 1990, or 40% of the total revenues counted above. In Office of Technology Assessment (1990), total U.S. energy use in 1985 is reported as 74.9 quads (quadrillion BTU's), of which 28 quads are reported for direct household energy use. This is 37% of the total. However, government direct use is also reported as 3 quads, so that uses in production (assuming that all energy use other than the two categories just mentioned should be counted

[4] From the Survey of Current Business
[5] From the 1990 Annual Energy Survey

as such) are only 59% of the total. Assuming that .6 of the costs calculated in the previous paragraph are energy inputs into non-energy production, we obtain energy costs with a value of 4.0% of GDP. Subtracting out the 5.5% of GDP representing value added by the domestic energy industry (6.6% minus 1.1% from above), value added in the non-energy sector represents 94.5% of GDP, so that energy costs in that sector are 4.2% of value added. The energy sector is thus not an extremely large one. It is thus somewhat surprising that taxes on the output of this sector have such large effects on aggregate activity.

2 Why Imperfect Competition Matters

We show below that the effects of energy taxes on aggregate activity are much larger when account is taken of imperfect competition among the firms that purchase energy. In this section we provide some intuition for this result by considering a simplified model. Suppose that output is produced with just two inputs, labor H and energy E. In particular, each firm has a Cobb-Douglas production function of the form

$$Y_t^i = A(H_t^i - \bar{H})^{1-\alpha}(E_t^i)^\alpha \tag{1}$$

where Y_t^i is the output of firm i in period t, while H_t^i and E_t^i represent its labor and energy inputs respectively. The parameter $\bar{H}$ represents a fixed amount of "overhead" labor needed to carry out any production at all. The assumption that there are fixed costs ensures that the production function exhibits increasing returns to scale in the sense that average costs exceed marginal costs. Our model requires us to assume such increasing returns to scale. Otherwise, it is impossible to reconcile the gap between price and marginal cost implied by the absence of perfect competition with the apparent absence of pure profits in U.S. industry.

Given the production function in 1, the marginal product of energy is $\alpha Y_t^i/E_t^i$ or, equivalently, $\alpha A[(H_t^i - \bar{H})/E_t^i]^{1-\alpha}$. Under perfect competition, this marginal product is set equal to the real price of energy, *i.e.*, to the price of energy divided by the price of output. But, under imperfect competition, the price of output is higher relative to marginal cost. In this case one instead obtains

$$\alpha A\left(\frac{H_t^i - \bar{H}}{E_t^i}\right)^{1-\alpha} = \mu_t^i p_{Et} \tag{2}$$

where μ_t^i is the ratio of firm i's price to its marginal cost in period t, and p_{Et} is the real price of energy at t. Equation 2 has two implications, both of which make energy taxes more contractionary in the case of imperfect competition. First, a high μ_t^i implies a higher marginal product of energy, given any observed real energy price. The fact that the marginal product of energy is higher implies that any given reduction in energy inputs lowers output by more under imperfect competition.

To see this more formally, note that 1 implies that a one percent reduction in E lowers output by α percent. The question is what value should be assigned to α. Under perfect competition, 2 implies that it equals the energy share $p_{Ei}E_i^i/Y_i^i$, and this is the usual method of assigning a numerical value to this parameter. But with a markup different from one, the energy share instead equals α/μ. Thus a higher markup implies a higher value for α, and thus a higher elasticity of output with respect to energy, given an observed energy share (as calculated in the previous section).

This still leaves the question of whether the energy input falls more under perfect or under imperfect competition. A second implication of 2 is that, holding employment fixed, the energy input falls more under imperfect competition. Holding employment fixed is reasonable if one expects labor to be supplied inelastically in the long run. Then 2 implies that a one percent increase in the price of energy will lead to a $1/(1-\alpha)$ percent reduction in the demand for energy. This fall is larger the larger is one's estimate of α, and thus the larger is the departure from perfect competition.

The intuition for this result is the following. Suppose that one observes that, with a given amount of employment, an economy produces 7 units of output with 50 units of energy input. Figure 1 displays two possible Cobb-Douglas production functions that could have led to this outcome. In the first α is equal to 0.5, while in the second α is equal to 0.7. They differ in that the marginal product of energy at the observed level of output is different. The function with α equal to 0.5 might be inferred, given the observed real price of energy, if one believed that firms are perfectly competitive, while the function with α equal to 0.7 might be inferred under imperfect competition. An important difference between the two functions is that the one with α equal to 0.7 is less bowed towards the origin, less concave. The smaller concavity of this function is dictated by the fact that both curves go through the origin and through point A, while the one with $\alpha = 0.7$ is steeper at A. The smaller concavity of the $\alpha = 0.7$ function implies that a given percentage change in its slope, *i.e.*, in the marginal product of energy, must lead to a larger change in the energy input. [6] Thus imperfect competition implies a larger change in the energy input from a given percentage tax on energy,

[6] The fact that the curve with α equal to 0.7 has both a steeper slope at point A and a flatter slope at low values of the energy input implies that the slope of this curve rises by less *in percentage terms* as one decreases the energy input from point A to a low positive value.

given observed values for output, the energy input, and the real price of energy at point A.

Under imperfect competition, the increase in energy taxes also has the potential of raising the equilibrium markup μ_t^i. It follows immediately from 2 that an increase in the markup will, with constant employment, lead to a further contraction in energy inputs and thus in output. Our simulations below show that in the case of a model of oligopolistic collusion, an increase in the energy tax does cause an increase in the equilibrium markup. In this case, imperfect competition has an even greater effect on our results.

3 A Simulation Model with Imperfectly Competitive Product Markets

As was noted above, our simulation model is similar in structure to the one used in Rotemberg and Woodford (1993) to analyze the effects of oil price shocks. Some modifications are required, however, for our present purposes. In particular, our interest in permanent tax changes requires that we take account of the effects of entry and exit in the long run. We also distinguish here between the use of energy in production and direct household use of energy.

The production function in our simulation model is much more general than the one used in the previous section for illustrative purposes. Like Goulder (1992), we assume that each firm in the private non-energy sector produces goods each period with a production function of the form

$$Y_t^i = Q(V(K_t^i, z_t H_t^i), G(E_t^i, M_t^i)) \tag{3}$$

where K_t^i, and M_t^i represent, respectively, firm i's capital and materials inputs at time t while z_t indicates an exogenously given labor-augmenting technology factor. The aggregator Q for value added V and the intermediate input aggregate G is assumed to exhibit constant returns to scale, as is the aggregator G for the intermediate inputs E and M. In the competitive case, we also follow Goulder in assuming constant returns to scale for the value added production function V. However, in the case of imperfect competition, and hence output prices higher than marginal cost in equilibrium, constant returns to scale would, again, imply the existence of pure profits. We do not wish to let such profits exist, at least not in the long-run steady-state growth path. Hence in the case of imperfect competition, we assume an increasing returns technology, so that average costs in excess of marginal costs can reconcile market power with free entry, as in Chamberlin's celebrated model of monopolistic competition. As in Rotemberg and Woodford (1992), we do this by assuming a value added production function of the form

$$V(K,H) = F(K,H) - \Phi \tag{4}$$

where F is homogeneous of degree one, and Φ is a positive constant. (We may assume that 4 applies equally in the competitive case, but with $\Phi = 0$.) The constant Φ indicates the presence of fixed costs (overhead), while the homogeneity of F implies that marginal costs are independent of scale.

We assume that z_t grows exogenously at a rate $g > 0$. The tax changes that we consider below will all be analyzed in terms of perturbations of the equilibrium around a steady-state balanced growth path that the economy would follow in the absence of the tax changes. Along this balanced-growth path, the aggregate capital stock, energy inputs, materials inputs, and non-energy output all grow at the same rate g (the exogenous rate of technical progress), while aggregate hours worked remain constant (so that the effective labor input $z_t H_t$ grows at the same rate as the other factors). [7] In order for fixed costs to remain a constant fraction of total costs along this balanced growth path, it is necessary for us to assume (in the case of imperfect competition) that the number of firms N_t grows at the same rate g, so that the scale of production by each firm remains constant. We assume that entry is through the introduction of new differentiated goods, so that the degree of market power of each firm remains the same (again, as in Chamberlin's model). The details of the process of entry and the conditions needed to ensure that our steady state with entry has zero profits are considered in Appendix 1.

We consider only symmetric equilibria in which the production plans of all firms are identical, so that $Y_t^i = Y_t/N_t$, $E_t^i = E_t/N_t$, and so on, where the variables without i superscripts refer to aggregate quantities for the private non-energy sector. The maximization of profits by these individual firms implies, as before, that the marginal product of each factor is equal to the product of this factor's real price and the markup of price over marginal cost. While there are four conditions of this type, we will mainly be interested in the one that is analogous to 2. This condition relates to the marginal product of G and requires that

$$Q_G(V_t, G_t) = \mu_t p_{Gt} \tag{5}$$

where p_{Gt} is a price index for the aggregate G_t, μ_t is the common markup of all firms in a symmetric equilibrium. The price index p_{Gt} depends on the prices of energy and materials relative to the price of

[7] In assuming a balanced growth path in which (per capita) hours worked remain constant, we follow numerous papers in the real business cycle literature; see, e.g., King, Plosser and Rebelo (1988). See also the footnote on page 10.

non-energy output. In a symmetric equilibrium the price of this output is the same for all firms, even in the case of imperfect competition. Because each firm's materials are some other non-energy firm's output, the price of materials inputs is identical to the price of non-energy output. Energy inputs are assumed to be in perfectly elastic supply at a fixed relative price p_E (which we imagine to be fixed on a world market, and so independent of changes in tax policy and production plans in the U.S.). Thus, p_{Gt} depends only on the tax rate on energy, τ_t, whose effects we wish to analyze. Because we assume that p_E is fixed in all of our experiments below, there is no distinction between the case of an *ad valorem* tax and a specific tax such as the BTU tax that was recently proposed.

In our simulations, we consider three different types of product market structure for the non-energy producers. In the case of perfect competition, equations 5 holds with $\mu_t^j = 1$ at all times. In our second model (the "constant markup" model), it holds with $\mu_t^j = \mu$, a constant greater than 1, at all times. This corresponds to a model in which firms are monopolistic competitors, with the equilibrium markup being determined by each firm's elasticity of demand, which in turn follows from the elasticity of substitution between the differentiated goods. [6]

Finally, in our third and most complicated model (the "variable markup" model), we assume that firms belong to oligopolies that maintain high prices through the threat of reversion to low prices if anyone deviates. Rotemberg and Woodford (1992) show that this implies that the markup μ_t^j for each firm in industry j will be related to the ratio of expected future profits to current sales. In particular, the markup will be given by

$$\mu_t^j = \mu(X_t^j / Y_t^j) \tag{6}$$

where $\mu(X/Y)$ is an increasing function, Y_t^j denotes the common output of each firm in the industry, and X_t^j denotes the expected present value of future profits gross of fixed costs for each firm in the industry assuming that collusion is maintained. Higher expected future profits relative to current sales raise the expected losses from a breakdown of collusion relative to the potential gains from undercutting the other firms in one's industry at the present time. The result is that collusion is easier to sustain. The formal definition of X_t^j can be found in Rotemberg and Woodford (1991) where we explain how X depends on the possibility that

[6] The assumption of a constant markup at all times does not actually require an assumption that the individual firm's demand curve has a constant-elasticity form, as in the familiar model of Dixit and Stiglitz (1977). Given that we consider only the symmetric equilibrium, it suffices that the utility received from the differentiated goods be a homothetic function, so that the elasticity of substitution between different goods *along the symmetric-consumption income expansion path* is constant. See Rotemberg and Woodford (1991) for further discussion of this model.

oligopolies will either be dissolved or renegotiate their collusive arrangements.

We now describe the rest of our simulation model. To model the supply of labor and capital, we assume the existence of a representative household that seeks to maximize

$$E\Big\{\sum_{t=0}^{\infty}\beta^t U(A(C_t, E_t^h), H_t^s)\Big\}$$

where β is a constant positive discount factor, C_t denotes consumption purchases of non-energy output (that for simplicity we here treat as entirely non-durable), E_t^h denotes household direct use of energy, and H_t^s denotes total hours worked (both for the private sector and for the government). The representative household is assumed to be a price-taker in all markets, and to face the wage w_t for all hours supplied, and the after-tax price of $p_E(1+\tau_t)$ for energy. (In some of our simulations below, we allow the tax on direct household energy use to differ from the tax on energy inputs to production.) The household also accumulates the capital stock (the purchase price of which is the same as the price of consumption goods), and receives the rental rate r_t on its capital holdings; and it owns all firms and receives the profits from both non-energy and energy production. Capital holdings evolve according to

$$K_{t+1} = I_t + (1-\delta)K_t$$

where I_t are period t investment purchases of non-energy output, and $0 < \delta \leq 1$ is a constant rate of depreciation.

In order to allow the existence of a balanced growth equilibrium in the case of a constant level of energy tax, we require as well certain homogeneity assumptions on household preferences. Specifically, we assume that the aggregator function $A(C, E)$ for household expenditure is homogeneous degree one. We also assume that the utility function $U(A, H)$ satisfies certain homogeneity assumptions explained further in Appendix II. These imply that if the household is faced with a real wage that grows at a constant rate and a constant rate of return on savings, it will choose to supply a constant number of hours, and to consume a quantity that grows in proportion to the real wage. [9]

[9] These assumptions are standard in the real business cycle literature. See, e.g., King, Plosser and Rebelo (1988). Apart from their analytical convenience, in allowing us to analyze a steady-state balanced growth path despite the existence of technical progress and endogenous labor supply, they are roughly accurate as a description of post-war U.S. growth. The most notable empirical embarrassment concerns not the growth of per capita private hours H_t, but per capita hours hired by the government, which exhibits a positive trend over the post-war period, contrary to the assumption of our model below. Needless to say, adequately dealing with the growth of the government sector observed over this period, if taken to represent a genuine long-run trend, would be incompatible with the existence of balanced growth.

As noted above, we assume that the supply of energy is infinitely elastic so that the relative price at which energy is supplied is fixed exogenously. This is probably not strictly correct. However, the view that the elasticity of supply is large is justified to some extent by the fact the price of oil is determined in a world market where the U.S. consumes only a quarter of world output. [10] Thus, even assuming that foreign demand is inelastic, the elasticity of supply faced by the U.S. is four times the world elasticity of supply. In addition, the foreign elasticity of demand also renders the effective supply of energy to the U.S. more price elastic. Put differently, any reduction in price brought about by a reduction in U.S. consumption would raise consumption elsewhere and thereby dampen the required fall in price. The result is that, even if the elasticity of the world supply of oil is zero, the effective elasticity of supply for the U.S. would equal three times the elasticity of demand of all the other nations. On the other hand, we abstract here from considerations of international trade by supposing that all U.S. energy usage (the sum $E_t + E_t^h$) is supplied by firms that are owned by the same representative household referred to above. [11] We also ignore for simplicity the use of factor inputs in energy production, and treat the revenues of the energy sector as pure rents (distributed as profits to the representative household).

We do take account of the consumption of real resources by the government, although in our simulations, government demand is assumed to simply grow deterministically with the rest of the economy. Specifically, we assume an exogenously given path for real government purchases of non-energy output $\{G_t\}$. In order to make possible a balanced growth path, we assume that G_t grows at the rate g of labor-augmenting technical progress. We similarly assume an exogenously given path for government purchases of people's time. In order to make possible a balanced-growth path of the kind described above, we assume that the hours per capita purchased by the government are a constant, H^g, at all times. We also assume that lump-sum taxes or transfers make up for any discrepancy in a given period between the value of government expenditure $G_t + w_t H^g$ and the value of energy tax revenues $\tau_t(E_t + E_t^h)$. This allows us to consider the effects of a change in the level of energy taxes while abstracting from the effects of changing other distorting taxes or of changing government expenditure patterns. Market clearing in the non-energy sector then requires that

[10] In 1989, the U.S, consumed 14.81 million barrels of oil a day while world production equaled 59.61 million barrels a day.

[11] We do assume in computing predicted changes in GDP that some of the energy is classified as "foreign" output for purposes of the national income accounts, but this is treated as an accounting convention with no economic significance. See equation 9 below.

at each time

$$C_t + I_t + G_t = Y_t - M_t \tag{7}$$

while market clearing in the labor market requires that

$$H_t + H_t^E = H_t^s \tag{8}$$

In our numerical simulations, we consider the comparative dynamics associated with deterministic perturbations of the expected time path of the energy tax $\{\tau_t\}$. In the case of perturbations that are small enough, the effects are essentially linear in the percentage tax change. The magnitude of these linear effects can be obtained from a log-linear approximation to the equilibrium conditions of the model. We carry out this linearization around the long-run steady-state balanced growth path to which the economy eventually converges. This allows us to state our results in terms of the percentage changes in non-energy output and so on *per percent increase in the energy tax*. It also means that the parameter values required in order to obtain numerical results are simply elasticities of the various functions introduced above, and average values of the various quantities. The parameter values used in our simulations are listed in Table 2. The sources of these numerical values, as well as the interpretation of the parameters, are discussed further in Appendix II.

In our basic simulation, we consider the effects of an permanent increase in the energy tax τ, that is announced (unexpectedly) at the same time that it takes effect. We assume that the economy had previously converged to the steady-state balanced growth path associated with the previous level of the energy tax (zero), and consider the path by which it converges to a new long-run steady state following the change. We also consider, for purposes of comparison, an experiment in which only the tax rate on direct household use of energy is increased, with no change in the tax on uses of energy as an input to non-energy production. In this case, the relative price of energy inputs in production continues to be p_E, while the relative price of energy for household use becomes $p_E(1+\tau)$. This comparison is of interest because the gasoline tax that was eventually passed as part of the 1994 Budget is, effectively, a tax that falls disproportionately on the energy purchased by households. It is thus of interest to compare the effect of such a tax to those of a more broad-based tax, such as the BTU tax originally proposed by President Clinton. As might be expected, we will show that imperfect competition increases the output losses associated with an energy tax *only* in the case of a tax on the use of energy in production. The reason is that imperfect competition affects the degree to which output falls only by affecting the energy purchases of firms.

4 Long-Run Effects of a Permanent Energy Tax

Table 3 summarizes the changes in the long-run levels of several variables, for each of four cases. The two types of tax changes considered are a shift from zero energy tax to a 1% tax on all energy use, (first two columns) [12] and a shift from no energy tax to a 1% tax on the direct use of energy by households (last two columns), assuming no tax on industrial uses of energy. Each tax change is considered for two alternative assumptions about product market structure. In the "competitive" case (left column of each pair), we assume perfect competition (*i.e.*, $\mu = 1$). In the "market power" case (right column of each pair), we assume imperfectly competitive product markets, with the typical firm possessing market power sufficient to lead it to set prices 20% higher than its marginal cost of production in the steady-state equilibrium (*i.e.*, $\mu = 1.2$).

As was noted above, our specification of a value for the steady-state markup μ also determines our specification of the degree of increasing returns in the production technology. In the "competitive" case, we assume constant returns to scale ($\Phi = 0$). In the "market power" case, we assume the existence of increasing returns due to the presence of fixed costs ($\Phi > 0$), and endogenous determination of the number of firms (and hence varieties of differentiated goods). Thus in this case there exist increasing returns such that average cost is 20% higher than marginal cost for the typical firm in the steady-state equilibrium. All other parameters are calibrated in the same way in the two cases.

One issue that arises at this point is whether a markup of 1.2 is reasonable. There are essentially two sources of information on this parameter. The first stems from the large literature which attempts to measure the elasticity of demand facing individual products produced by particular firms. This literature is relevant because it is never profit maximizing for a firm to set its markup lower than one over one plus the inverse of the elasticity of demand for its product. There are many estimates of the elasticity of demand for particular products in the marketing literature. Tellis (1988) surveys this literature, and reports that the median measured price elasticity is just under 2. Thus the markup would equal at least 2 if this sample of firms is representative. In practice, elasticities of demand undoubtedly differ across products and the elasticity of demand of those products studied in the marketing literature is probably atypically low. This is because the

[12] Although we assume here an initial steady state with no energy tax, the results would be similar in the case of a 1% increase in the value of $(1+\tau)$, starting from a positive initial tax.

marketing literature focuses on the demand for branded consumer products which are more differentiated than unbranded products so that their demand is probably less price sensitive. Thus, the typical product in the economy probably has a price elasticity of demand that exceeds 2.

A second approach is to analyze what happens to revenue and costs in response to an exogenous change in aggregate demand. A particularly simple version of this approach has been proposed by Hall (1988, 1990). He studies the degree to which the increase in GDP generated by increases in exogenous variables such as changes in military purchases is accompanied by an increase in costs. Insofar GDP increases by more than costs, the markup is greater than one. His estimates indicate that the markup μ is between 1.4 and 1.6. [13]

There is also a related literature which tries to obtain econometric estimates of marginal cost and, in some cases, combine them with econometric estimates of the elasticity of demand. The aim of this approach is to obtain simultaneous, independent estimates of the markup and of the degree of increasing returns. Morrison (1990), for example, estimates a flexible functional form cost function, using data on gross industry output and materials inputs. Her estimates of that markup μ range between 1.2 and 1.4 for 16 out of her 18 industries. One notable feature of these estimates is that her industry estimates of the ratio of average to marginal cost closely resemble her estimates of the markup itself. Thus the relation between these two parameters that we imposed through our zero profit condition appears to be validated.

Because we are considering only long-run effects, the results do not depend on whether the tax increase is immediate or phased in over a period of time; only the eventual permanent increase in the tax rate matters. Similarly the results do not hinge on whether the long-run substitutability of factors of production exceeds their short-run substitutability; only the long-run substitution possibilities matter here.

Furthermore, the "market power" case reported in Table 3 refers equally to the monopolistically competitive model and to the oligopolistic collusion model. The reason is that in neither case does the energy tax change have any effect on the markup of prices over marginal cost in the long-run steady state. In the case of monopolistic competition, the markup is predicted to be a constant, determined solely by the elasticity of substitution between alternative differentiated goods. In the oligopolistic model, by contrast, the markup depends upon the ratio X/Y, and so can vary in response to policy changes. However, in a

[13] His reported estimates for markups are actually even higher. The reason is that he estimates "value-added" markups as opposed to the more standard markups of price over total marginal cost. For a discussion of the relation between the two, see Rotemberg and Woodford (1992).

steady-state equilibrium, the present discounted value of profits X is proportional to Y. Moreover, while the steady state value of X also depends on the steady-state real rate of interest, r, this real rate of interest is solely determined by preference parameters and the exogenous rate of growth (see Appendix II). Hence the steady-state r is unaffected by the energy tax, as a result of which steady-state μ is also unaffected. Thus the long-run effects are the same in either type of imperfectly competitive model; all that matters is the size of the steady-state markup μ.

We now turn to the numerical results reported in Table 3. In each row, the figure reported represents the *percentage change* in the long-run value of that variable resulting from a 1% energy tax. (In our log-linear approximation to the equilibrium, the effects of a k% energy tax are obtained by multiplying each of these numbers by k.) The variables, of course, grow over time in the steady-state equilibrium; but the steady-state growth rate is unaffected by the energy tax (as it is determined solely by the exogenous rate of technical progress). Thus the figure -.071 for non-energy output means that output is -.071 percent lower at all times than it otherwise would have been, in the new long-run steady-state growth path. "Non-energy output" refers to the gross output Y_t of the energy-using (but not energy-producing) sector. The change is "GDP" is computed as

$$\Delta Y_t - \Delta M_t - p_E \Delta E_t + \theta^{dom} p_E(\Delta E_t + \Delta E_t^h) \tag{9}$$

where Δ indicates the difference (*not* percentage difference) in value of the equilibria between the perturbed and unperturbed equilibria, and θ^{dom} denotes the share of energy used in the U.S. that is domestically produced. [14] Thus the GDP measure aggregates value added in the non-energy sector and the domestic energy-producing sector, where for simplicity the total revenues of the latter sector are counted as value added. "Hours worked" denotes total hours worked H_t; because government hours are assumed to follow an exogenous path unaffected by the energy tax, the reported decline in hours is only 83% of the size of the decline in hours in the private non-energy sector. The "product wage" refers to the wage deflated by the price of non-energy input (*i.e.*, , the quantity w_t in equation (3a)), while the "consumption wage" refers to the wage deflated instead by the price index p_{At} of the household consumption basket,

$$p_{At} \equiv \frac{C_t + p_E(1+\tau_t)E_t^h}{A(C_t, E_t^h)} \tag{10}$$

[14] Implicitly, we assume here that U.S. energy production varies in the same proportion as U.S. energy use.

"Energy use in production", "capital stock" and "number of firms" refer to the variables E_t, K_t, and N_t introduced in the previous section; all refer solely to the private non-energy sector. The number of firms is indeterminate in the competitive model.

A striking feature of the results in Table 3 is that a tax that is levied only on direct household purchases of energy has *no effect whatsoever* on equilibrium activity in the private non-energy sector. Household energy use falls, and the consumption wage falls because the price index p_A rises. However, the household does not change its supply of labor or demand for non-energy goods, nor does the equilibrium product wage in the non-energy sector change. GDP falls only because of the reduction in domestic energy production due to reduced household use of energy. [15]

Two features of our model account for this result. The first is that we made assumptions that ensured the existence of a steady state where the economy grows but hours worked do not. Since output, consumption, energy purchases and wages all grow at the same rate in such a steady state, we require that equiproportional increases in wages and the aggregate $A_t = A(C_t, E_t^h)$ be consistent with an unchanged quantity of labor supplied. [16] A permanent increase in the tax on household energy raises $p_E(1+\tau_t)$ and thus raises the consumption deflator p_{At} while it lowers the consumption wage. As long as A_t falls in exact proportion to the increase in p_A while the product wage is unaffected, the fall in the real wage and in consumption are equiproportional so the quantity of labor supplied does not change.

The second important source of this result is our assumption that the household's elasticity of substitution between non-energy consumption and direct energy use is equal to 1. We base this on the unit elasticity of demand estimates of Houthakker, Verleger, and Sheehan (1974). This means that the shares of household expenditure on energy and non-energy output will remain constant in the face of a change in the relative price of these two kinds of goods. But supposing that the path of $p_{At}A_t$ is unchanged under the circumstances just described, it follows that the paths of C_t and of $p_E(1+\tau_t)E_t^h$ are unchanged. Thus household energy use falls in inverse proportion to the tax increase, while non-energy consumption demand and labor supply are unchanged. This means that, if the product wage and the real rate of return are unaffected, consumption demand and labor supply are unchanged so that output in the non-energy sector can remain constant as

[15] Goulder (1993b) also reports smaller losses in GDP for taxes that affect only household's use of energy.

[16] To obtain this result we assume that the Frisch consumption demand curve and the Frisch labor supply curve satisfy certain homogeneity properties explained in Appendix II.

well. This in turn implies that the previous paths of the product wage and the real rate of interest continue to describe an equilibrium. This argument applies whether there is perfect competition or not since the two models differ only in the way non-energy firms react to changes in their environment. But, as we just saw a tax on household use of energy does not affect this environment.

Since the tax on household direct use of energy has no effect on non-energy output, the effects on output of a tax on all energy are due to the tax on the industrial use of energy. The effects would have been just as large if *only* the energy used in production were taxed. It is true that the result of exactly zero effect of the tax on direct household use depends upon particular parameter choices that might well be challenged. However, for any values near ours, the result will still be approximately true – the effects of a tax on household energy use will be much smaller than the effects of a tax on industrial uses, and indeed the effects of a tax on household energy use could as easily be expansionary as contractionary. Thus the shift from a "BTU tax" to a gasoline tax in the budget that was eventually passed by the U.S. Congress probably resulted in a tax that places less of a burden on the economy, per dollar of revenue raised. Another implication is that in designing a tax intended to reduce carbon dioxide emissions, a tax aimed more at household energy use is likely to contract economic activity less, for any given reduction in emissions that is achieved, than one based simply on the "carbon content" of various fuels.

Now consider the effects of a tax on all types of energy use. Private non-energy production contracts, as do hours worked, the energy used in production, the capital stock and the real wage deflated by the price of non-energy output. The contraction in hours stands in contrast to the case of a tax on energy consumption only. It comes about because a tax on all energy lowers the consumption wage by more than it lowers A_t. The table shows that GDP falls by slightly more than does non-energy output in this case. The reason is that the contraction of the energy sector is even more severe, in percentage terms, than the contraction of the private non-energy sector. In the imperfectly competitive case, there is also a reduction in the long-run number of firms, due to exit in response to profits no longer large enough to cover the fixed costs.

Even in the competitive case, the output lost as a result of the energy tax is rather significant. Since the share of total energy expenditure in GDP is 0.066, a one percent tax increase raises government revenues by only .066 percent of GDP. On the other hand, GDP is itself reduced (in the long run) by .071 percent.

[17] The ratio of output loss to revenue raised is even more severe if one considers the case of a pure tax on industrial uses of energy. In this case, a one percent energy tax raises government revenues by only .040 percent of GDP, while GDP is reduced by .050 percent. [18]

But the contractionary effects of an energy tax are even greater when one allows for imperfect competition. In the case of market power, the long-run decline in non-energy output is .097 percent, and the long-run decline in GDP is .098 percent. Thus even a relatively modest degree of market power (prices 20% above marginal cost) significantly increases the predicted effect of the energy tax; the long-run non-energy output decline is increased by a factor of nearly 1.4.

The reason for this can be understood by analyzing equation 5 in the same way that we analyzed equation 2 in section 2. It follows from 5 that a value $\mu > 1$ implies a higher estimate of the elasticity of the aggregator Q with respect to G. Using the shares reported in Table 2, the implied value of this elasticity is increased from .52 to .624 (a factor of 1.2) by raising μ to 1.2. The elasticity of Q with respect to V similarly falls from .48 to .376. Furthermore, as in section 2 the higher value of μ implies that the fall in G is larger in response to a change in the price p_G. It turns out that, for fixed V, the fall in G is inversely proportional to the elasticity of Q with respect to V, and thus is larger by a factor of 1.28 when $\mu = 1.2$. Thus, supposing that the index of primary inputs V were not affected by the energy tax, G would have to fall by 1.28 times as much, and so output would fall by 1.53 (= 1.2×1.28) times as much, in the case of market power. In fact, the difference between the output declines in the two cases is not so extreme in our simulations. This is because V actually falls more in the long run in the competitive case. Nonetheless, the contraction is significantly larger in the presence of market power.

The energy tax also has significantly greater adverse effects in the case of market power in several other respects as well. For example, the real product wage falls by more than 1.5 times as much in percentage

[17] This ratio of the output loss to the revenue raised is comparable to what is implied by the results of other authors. For example, Goulder (1992) estimates that a $25 (in 1990 prices) per ton tax on carbon content will reduce real GNP in 2020 by an amount that varies between .76 percent and 1.14 percent, depending upon the structure of the tax (see his Table 7). Using the figures in his Table 1 on the percentage tax that this corresponds to for different types of fuel, one finds that this corresponds to an average tax rate on energy use of 14.5%, so that the revenues raised should be only approximately .96 percent of GNP, even ignoring the reductions in the cost share of energy that should follow from such a severe tax. Goulder (1993b), on the other hand, reports smaller GNP losses. He considers a tax of 45 cents per million BTU's which corresponds roughly to a 22% tax on energy and computes a reduction in GNP of only a third of one percent. By contrast our estimates would imply that a 22% tax would lead to GDP losses of 1.5 percent.

[18] Non-energy output still falls by .071 percent in this case, but the energy sector contracts to a much smaller extent. Total energy sales are affected more by a tax on household energy use than by a tax on industrial uses, because our calibration implies that energy is more substitutable for other goods for households than for firms.

terms. This indicates a more significant contraction of labor demand; the only reason that hours worked do not fall more is that households are willing to accept the lower real wage because of their lower wealth in this case.

5 Short-Run Effects of Changes in Tax Policy

In this section we begin the more complex analysis of the short-run effects of changes in policy. We focus here on the transition to a new long-run steady state consistent with a new permanent tax rate. During this transition the effects of tax rate changes differ depending on whether markups are constant or not. We thus consider separately our two models of imperfect competition.

We consider here an unexpected permanent increase in the energy tax rate τ from zero to .01 that takes effect immediately. Since the tax change is not anticipated in advance, we suppose the economy starts out in the steady-state balanced growth path associated with zero energy taxation. We imagine that the tax applies to all uses of energy, although as explained in the previous section, the effects of the tax on the non-energy sector follow solely from the taxation of energy used as an input to non-energy production. [19]

Figure 2 displays the transitional effects on non-energy output, under the three alternative specifications of product market structure. The vertical axis indicates percentage deviations from the previous steady-state growth path; -10×10^{-4} means a reduction of .10 percent. The horizontal axis indicates the year; year 0 is the year in which the tax change is announced and takes effect. In the competitive case, the tax lowers non-energy output by .071 percent in the long-run as we showed in Table 3. We now see the short-run effects as well. In the first year, non-energy output is already reduced by .058 percent. In subsequent years, output continues to fall further below the previous trend path, as the capital stock is eroded; but a large part of the eventual output decline occurs immediately. In the case of the constant markup model, the general picture is similar. But, as the elasticity of output with respect to energy inputs is larger in this case, the decline in output is larger both in the short run and in the long run.

In the case of the variable markup model, the long run effects are the same as for the constant markup model. As explained in the previous section, both models predict a long-run reduction of non-energy output

[19] The argument made above shows that in the case of an immediate, permanent tax on direct household energy use only, the economy will immediately jump to a new steady-state balanced growth path described in Table 3. Thus there is no effect upon non-energy output, consumption, and so on, in either the short run or the long run.

by .097 percent. However, the short run effects are quite different. During the first year of the energy tax the constant markup model predicts a reduction by .083 percent, while the variable markup model predicts a reduction by .123 percent. The predicted short run effect is almost one and one-half times as large (and thus is more than twice the size of the effect in the competitive model). This is because the variable markup model predicts that the markup increases when the energy tax is increased, and then gradually returns to its original level over time. The markup increases because the ratio X/Y increases. This occurs, in the first instance, because of a decline in real interest rates that results from the reduced returns to capital (which eventually return to normal as the capital stock is reduced). Lower real interest rates mean that the expected future profits from collusion are discounted to a lesser extent, making a greater degree of collusion possible. Higher markups then themselves contribute to a higher ratio of profits to sales each period, making X/Y still higher and so helping to raise markups further. Higher markups also further reduce the returns to existing capital goods. Thus lowers real rates of return further, thus further raising X/Y and further raising markups in a self-reinforcing process.

In our simulation, the markup increases by .011 percent during the first year of the tax (i.e., from 1.2000 to 1.2001). Even this small increase in the inefficiency wedge due to firms' market power has a significant effect on the predicted equilibrium allocation of resources. To understand this, it is helpful to suppose first that labor supply is inelastic. Then, V_t is entirely As a result, 5 determines G_t as a function of (τ_t, μ_t). Then, using 3, Y_t depends only on (τ_t, μ_t). Now let us investigate for a given increase in τ_t, the quantitative effect of a an increase in μ_t. Because energy costs are only about 4% of total intermediate input costs, a one percent increase in the after-tax energy price raises the price index p_{Gt} by only .04 percent. Thus a contemporaneous .011 percent increase in the markup means that the right hand side of 5 increases by 1.3 times as much (in percentage terms) as it would in the case of a constant markup. In our log-linear approximation to the solution, the percentage decline in Y_t is proportional to the percentage increase in the right hand side of 5, and so it should be 1.3 times as large in the variable markup case.

In our simulation model, we also allow for endogenous labor supply. [20] In this case, households reduce labor supply rather than accept a real wage cut of the size that would be required to induce firms not to

[20] Note that it is not essential to our conclusions here that this be interpreted as "voluntary" variation in labor supply. Qualitatively similar conclusions would be obtained in the case of any source of short-run "real wage rigidity", due for example to pre-existing wage contracts or to efficiency wage considerations.

reduce the labor inputs that they use. Thus non-energy output falls even more than it would in the case of inelastic labor supply. This effect is present regardless of product market structure. However, one can easily see that the real wage decline required to induce firms not to reduce labor inputs is larger if the markup rises. For it follows from 5 that if the markup rises, the value of G_t falls more, and hence that $Q_V(V_t, G_t)$ for fixed labor and capital inputs falls more. On the other hand; the same logic that leads to 5 implies that the real wage must equal $Q_V V_H/\mu$. The fall in the real wage is thus magnified both by the increase in μ and by the severity in the fall of Q_V. It thus makes sense that in the case of endogenous labor supply, the effect of the markup increase on output is even greater. It is the fact that such small changes in the markup can matter so much for the size of the predicted effects of the tax increase that leads us to insist upon the importance of product market structure for tax analyses of this kind.

6 The Effect of Expected Changes in Energy Taxes

Up to this point we have considered the effect of unanticipated permanent increases in energy taxes. There are several reasons, however, for energy taxes to be anticipated. First, there is a time gap between the moment where tax policy is announced and when it takes effect. In particular, the Clinton proposal called for a gradual phase-in of the BTU tax. Second, tax changes are not necessarily permanent. Any particular tax, such as the energy tax has some probability of being repealed in the future.

We start by considering the case of gradual phase-in. We thus report simulations in which the energy tax is increased by one-half of one percent in the year that it is announced while the full one percent tax applies from the second year onward. The comparison with the case of an instantaneous increase in the tax is interesting in part because a gradual phase-in was actually proposed. Moreover, as we will show, the effect of this gradual phase-in depends even more crucially on product market structure than the eventual effect of a permanent tax.

The consequences for non-energy output are shown in Figure 3, for each of the possible market structures. As one might expect, output falls by less in the first year than if the full tax were to take immediate effect. In fact, in none of the models is the contraction in the first year even half the size indicated in Figure 1. In the case of the competitive model, the first-year decline in non-energy output is only .016 percent; in the constant markup model, it is only .012 percent; and in the variable markup model, output does not decline at all in the first year, but instead *rises* by .011 percent.

One reason for the first year effect to be so muted is that the wealth effect on labor supply, which tends to increase equilibrium output, is nearly as large in these simulations as in the previous ones. On the other hand, the current increase in energy costs, which tends to reduce equilibrium output, is only half as large. The other important factor, in the case of the variable markup model, is that expectations of future profits are reduced nearly as much as in the previous simulations, while current sales are reduced by much less. This means that the ratio X/Y falls, so that the equilibrium markup is temporarily reduced in the oligopolistic model. The path of the equilibrium markup in the oligopolistic model is shown in Figure 4, for the cases of the immediately effective tax and the phase-in over a one-year period. In the case of the gradual phase-in, the markup falls by about .004 percent (*i.e.*, , from 1.2 to 1.19995).

After the first year, the path of output in these simulations is similar to the one we derived for an immediate tax increase. The only difference is that the higher output in year 0 is associated with a higher level of investment. Thus the capital stock in year 1 is higher. In fact the economy now begins year 1 with a slightly larger capital stock than in the original balanced-growth path in all cases. By contrast, the capital stock was slightly lower in each of our previous simulations. In the case of the competitive model and the constant markup model, the higher capital stock means that the output decline in year 1 and later is not quite as large as in the previous simulations. On the other hand, in the case of the variable markup model, we find that a higher capital stock actually makes the output decline even more severe. The higher capital stock implies that real interest rates are even lower. This implies that X/Y is even higher and thus leads to even higher markups. Figure 4 shows that, indeed, the markups in year 1 and later are actually greater in the case of a phased-in tax.

Finally, we report simulations in which the tax increase is not expected to permanent. We now suppose that the tax is increased to 1% on all uses of energy, but that it is anticipated that each year there is a 20% probability that the tax rate will be permanently restored to its original value. In our dynamic equilibrium model,the effects of a tax increase cannot be analyzed independently of expectations about future policy, and it is important to realize that economic agents need not expect that a tax change is permanent simply because the bill that is enacted does not specify a future date at which it becomes invalid. Here too we find that the effects of an expectation of future policy reversal depend greatly upon our assumptions about product market structure.

Figure 5 presents the time path of non-energy output in the case of the three possible market structures. Here what is plotted for each year is the level of non-energy output in that year relative to the previous trend growth path assuming that there has been no reversal of the tax up until that time. In the case of both the competitive model and the constant markup model, the contraction of non-energy output is greater than it would have been were the tax expected to continue forever. This is due to the wealth effect on labor supply; optimism about reversal of the tax makes households expect higher future incomes and thus makes them less willing to work in the present. This demonstrates that the contractionary effects of energy taxes may, in practice, be considerably greater than those indicated in Figure 1. [21]

In the case of the variable markup model, things are more complex. It is again true that the expectation of a possible reversal lowers first period output because of the wealth effect on labor supply. However, the possibility of a policy reversal also raises the equilibrium real rate of return, because higher rental rates on capital are expected in the event of repeal of the tax. This higher rate of return lowers the present discounted value of future profits relative to current revenues. The resulting reduction in X/Y lowers equilibrium markup in the oligopolistic model. While the markup still rises following enactment of the tax, it does not rise as much as in the simulation depicted in Figure 1. And, assuming that the tax has not yet been repealed, the equilibrium markup from year 3 onward in the oligopolistic model is actually lower than that in the monopolistically competitive model. This occurs because, once the capital stock has fallen sufficiently below its initial level, the real rate of return remains consistently above the real rate associated with the initial steady state. The consequence is that if, contrary to expectation, the tax continues for many years, output is actually higher in the oligopolistic model than in the monopolistically competitive model.

7 Conclusions

We have found that allowing for imperfect competition in product markets has an important quantitative effect on estimates of the effects of energy taxes on the level of economic activity. Allowing for even a modest average markup of prices over marginal cost increases the predicted decline in output which is caused by an

[21] There are other reasons why one might expect the stimulus to labor supply from the expectation of low future incomes not to be as large as in the simulations depicted in Figure 1. For example, one might suppose that some suppliers of labor are unable to borrow against future income in any event. In such a case, the contractionary effects of an energy tax are likely to be larger than those indicated in those simulations.

increase in the after-tax relative price of energy inputs. And allowing for even a small increase in equilibrium markups, due to increased sustainability of collusion among members of an oligopoly, can greatly increase the predicted output decline.

We have paid particular attention to a specific model of oligopolistic collusion that we have elsewhere argued helps to explain the responses of the U.S. economy to a variety of kinds of macroeconomic shocks. This model implies that an increase in energy taxes may well temporarily raise equilibrium markups., especially when it is both unexpected and expected to be reversed soon with significant probability. In this case, the short-run contractionary impact of an energy tax is especially large. This effect is, however, sensitive to the precise dynamic specification of the proposed taxes. Markups in the oligopolistic model may fall rather than rising immediately following announcement of an energy tax increase if there is a delay in the implementation of the tax.

In general, our results suggest an even less favorable relation between the revenues raised by an energy tax and the reduction of economic activity than earlier studies (assuming competitive markets) have indicated. For example, in the case of immediate implementation of a 1% energy tax that is expected to be reversed each year with a 20% probability, the revenues raised in the first year of the tax will be .066 percent of GDP, while GDP is itself reduced by .110 percent in the first year according to the constant markup model, and by .142 percent according to the variable markup model. The GDP reduction five years later is only .098 percent in the variable markup model if, contrary to expectation, the tax increase has not yet been reversed; but it is by that time .134 percent in the constant markup model. Although we do not here analyze alternative revenue sources, we believe that an energy tax is relatively unattractive on this dimension.

Our results also suggest ways in which an energy tax might be structured to minimize the contractionary effects. Our most important finding in this respect is that a tax solely on direct household use of energy need not contract non-energy production at all. Insofar as allowing for imperfect competition increases the predicted contractionary effects of a tax on industrial uses of energy, but does not affect the predicted consequences of a tax on direct household use, it makes the case for targeting household energy use even stronger.

The short-run contractionary impact of an energy tax is also reduced if the tax is phased in gradually, and our simulations indicate that the output gained in the transition period is much larger than the revenue losses due to the gradual phase-in. In the case that we analyze here, for example, gradual phase-in involves a revenue loss of .033 percent of GDP in the first year relative to the revenues from immediate implementation. But the result is that GDP falls by .070 percent less in the case of constant markups equal to 1.2, and by .109 percent less in the case of the variable markup model. In the case of the constant markup model, the output loss is also somewhat mitigated in later years although it is made slightly worse in the case of the variable markup model.

8 References

Diamond, Peter A. and James A. Mirrlees, "Optimal Taxation and Public Production I: Production Efficiency and II: Tax Rules", *American Economic Review*, 61, 1971, 8-27 and 261-78.

Dixit, Avinash, and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67, 1977, 297-308.

Goulder, Lawrence H.,"Carbon Tax Design and U.S. Industry Performance," *Tax Policy and the Economy*, 1992, 59-104.

———,"Effects of Carbon Taxes in an Economy with prior Tax Distortions: An Intertemporal General Equilibrium Analysis," mimeo, 1993a.

———,"Energy Taxes: Traditional Efficiency Effects and Environmental Implications," *Tax Policy and the Economy*, this issue. Other Taxes are Accounted for?" *NBER Working Paper no. 4061*, May 1992b.

Hall, Robert E.:"The Relation Between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, 96, October 1988, 921-48.

—:"Invariance Properties of Solow's Productivity Residual," in Peter A. Diamond ed.Growth, Productivity and Unemployment, Essays to Celebrate Bob Solow's Birthday, Cambridge: MIT Press, 1990.

Houthakker, H.S., Philip K. Verleger, Jr., and Dennis Sheehan,"Dynamic Demand Analyses for Gasoline and Residential Electricity," *American Journal of Aricultural Economics*, May 1974, 412-18.

Jorgenson, Dale W., and Peter J. Wilcoxen, "Reducing U.S. Carbon Emissions: An Econometric General Equilibrium Assessment," *Resource and Energy Economics*, 15, March 1993, 7-25.

Judd, Kenneth L.,"The Optimal Tax on Capital Income is Negative," mimeo, 1993.

King, Robert G., Charles I.Plosser, and Sergio Rebelo,"Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics*, 21, March 1988, 195-232.

Morrison, Catherine J.:"Market Power, Economic Profitability and Productivity Growth Measurement: An Integrated Structural Approach," *NBER Working Paper 3355*, May 1990.

Office of Technology Assessment, Congress of the United States, *Energy Use and the U.S. Economy*, U.S. Government Printing Office, June 1990.

Pindyck, Robert S., and Julio J. Rotemberg, "Dynamic Factor Demands and the Effects of Energy Price

Shocks," *American Economic Review*, 73, 1983, 1066-1079.

Rotemberg, Julio, and Michael Woodford, "Markups and the Business Cycle," *NBER Macroeconomics Annual*, 1991, 63-129.

———,"Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, 100, 1992, 1153-1207.

———,"Imperfect Competition and the Effects of Energy Tax Increases on Economic Activity," mimeo, April 1993.

Tellis, Gerald J.:"The Price Elasticity of Selective Demand: A Meta-Analysis of Econometric Models of Sales," *Journal of Marketing Research*, 25, November 1988, 331-41.

Appendix I

Entry and the Elimination of Profits

We have stated above that in the long run, entry and exit are assumed to maintain pure profits at zero. It is straightforward to show that the first order conditions for profit maximization imply that pure profits are zero in a symmetric equilibrium if and only if

$$(\mu_t - 1)Y_t = Q_V(V_t, G(E_t, M_t))N_t\Phi \tag{11}$$

where μ_t denotes the common markup of all firms. Equation 11 would thus determine the equilibrium number of firms each period in the case of instantaneous entry and exit. This equation refers only to the case of imperfect competition and increasing returns where μ exceeds one and Φ is strictly positive. Otherwise, as usual with constant returns, the number of firms is indeterminate.

We do not, however, suppose that entry and exit occur so quickly. Because entry and exit are peripheral to our main interests here (and because, as long as they are slow, the exact dynamics do not matter much for our results) we adopt a simple *ad hoc* specification rather than explicitly modeling the entry and exit decisions. Let us define

$$\bar{N}_t \equiv \Phi^{-1} lim_{k\to\infty}(1+g)^{-k} E_t\left\{\frac{(\mu_{t+k}-1)Y_{t+k}}{Q_V(V_{t+k}, G(E_{t+k}, M_{t+k}))}\right\} \tag{12}$$

Thus $\bar{N}_t$ denotes the number of firms needed at date t, if a constant rate of growth g of the number of firms ever after is to result in zero profits in the long run. We then assume that the number of firms grows exogenously at the rate g, except for a slow tendency to correct any discrepancy between the current number of firms and $\bar{N}_t$. Specifically, we assume dynamics for the number of firms given by

$$N_t = \rho\bar{N}_t + (1-\rho)(1+g)N_{t-1} \tag{13}$$

where $0 < \rho \leq 1$ is a constant partial-adjustment rate. This specification introduces an additional predetermined state variable, in addition to the aggregate capital stock K_t, and that is the previous number of firms N_{t-1}. Note that once there ceases to be new information about future policy, $\bar{N}_t$ grows at the constant rate g, so that 13 implies that the percentage discrepancy between N_t and $\bar{N}_t$ is eliminated at an exponential rate. Substitution of 12 and comparison with 11 indicates then that the share of pure profits in total revenues must asymptotically approach zero, as desired.

Appendix II

Parameter Values Used in Simulations

Here we explain the numerical values reported in Table 1. The steady-state balanced growth path of the economy is described by a set of growth rates and shares, that we calibrate using the U.S. national income accounts. According to our model, the exogenous rate g of labor-augmenting technical progress is also the steady-state rate of growth of real GDP, which is why we assign the value .03/year. The parameter r represents the real rate of return in the steady-state equilibrium. This is not a primitive of the model, but the model predicts that it should equal $\beta^{-1}(1+g)^{\sigma}$, so that calibration of r is equivalent to calibration of the rate of time preference of the representative household. Following King, Plosser and Rebelo (1988), we calibrate r to match the average real return on the U.S. stock market. The parameter δ represents the exogenous rate of depreciation of the capital stock of the private non-energy sector. The model implies that in a steady-state equilibrium, the share of investment in final uses and the share of capital in total costs must be linked, through the relation

$$\frac{s_I}{g+\delta} = \frac{s_K}{(r+\delta)(1-s_M)}$$

Hence the values assumed for g and r, and the share parameters discussed below, imply a value for δ, which is the one given. These parameters imply a steady-state capital-output ratio in the private non-energy sector of 7.5 quarters, which is reasonably consistent with the national income accounts as well.

The parameters s_C, s_I, s_G represent the steady-state shares of private consumption expenditure, private investment expenditure, and government purchases of private non-energy output, respectively, in total final uses $Y - M$ of private non-energy output. We calibrate these shares to equal the average shares of these three kinds of expenditure in U.S. private value added (GDP minus value added by the federal, state and local governments). The parameters s_E, s_M, s_H, s_K represent the steady-state shares of energy, materials, labor and capital costs, respectively, in the total costs of the private non-energy sector.

As we explained in section 1, energy costs in the non-energy sector are 4.2% of value added. Hence we must have

$$\frac{s_E}{1-s_E-s_M} = .042$$

We assume, somewhat arbitrarily, a share of materials costs of .5. This is somewhat smaller than the average materials share indicated in the Commerce Department data for U.S. manufacturing sectors, but we suppose that materials are a smaller fraction of costs outside of manufacturing. The above equation then implies $s_E = .02$. This leaves .48 of total costs for labor and capital costs. Insofar as wages account for about 75% of value added in the national income accounts, we set $s_H = .36$, $s_K = .12$.

As is explained in the text, we assume that in the long run, the number of firms is such that equation 11 is satisfied. This implies that in the steady state, s_Φ, the share of fixed costs in total costs, must equal

$$(1 - s_E - s_M)\frac{N\Phi}{F(K, zH)} = \frac{\mu - 1}{\mu}$$

Hence our calibration of this parameter follows from our choice of μ, discussed below. (Note that s_Φ does not refer to costs in addition to the four categories previously listed. The fixed costs are a subset of the costs already counted once as labor and capital costs.)

The calculations just explained imply that the share θ^{dom} of total energy use that is domestically produced is .83. They also explain why we set θ^h, the share of total energy use that is direct household use, equal to .4. The parameter η^g, indicating the steady-state value of H^g/H^p, is set equal to .17, the average ratio of government employment (summing employment by federal, state and local governments) to total employment over the postwar period.

This completes our specification of the parameters describing balanced growth. We turn next to the remaining parameters of the production technology of the private non-energy sector. As the functions Q, F, G are all assumed to be homogeneous of degree one, the only parameters that occur in the log-linear approximation to our equilibrium conditions are, in the case of each function, the elasticity of the function with respect to each factor (only one free parameter per function as they must sum to 1) and the elasticity of substitution between the two factors. The elasticities of substitution enter the log-linear approximations to those equilibrium conditions involving marginal products. All of these elasticities are evaluated at the factor mix that occurs in the steady-state equilibrium. The elasticities with respect to the individual factors are implied by the steady-state share parameters already discussed; for example, the elasticity of G with respect to E must equal $s_E/(s_E + s_M)$, and the elasticity of Q with respect to V must equal $1 - \mu(s_E + s_M)$. (It will be observed that for both of the values of μ that we use, each of these elasticities is positive.) It thus remains only to specify the elasticities of substitution. The values given in Table 1 for ϵ_{GV} and ϵ_{EM} are based upon the econometric estimates reported in the Appendix of Rotemberg and Woodford (1993). The value of 1 for

ϵ_{KH} (which would follow from a Cobb-Douglas production function for value added) is standard in the real business cycle literature and in a great many other computational general equilibrium studies.

We next consider the parameters of household preferences. As noted above, the rate of time preference is implicitly determined by our specification of r. It is useful to discuss the utility function $U(A, H)$ in terms of the Frisch demand functions $A^d(w^A, \lambda), H^s(w^A, \lambda)$ that it implies, where w^A denotes the "consumption wage" defined in section 2, and λ denotes the representative household's marginal utility of wealth (with wealth in units of the composite good A). [22] In order for a steady-state balanced growth path to be possible, it is necessary to make a homogeneity assumption on the Frisch demands. [23] Specifically, we assume that there exists a $\sigma > 0$ such that $H^s(w^A, \lambda)$ is homogeneous of degree zero in $(w^A, \lambda^{-1/\sigma})$, while $A^d(w^A, \lambda)$ is homogeneous of degree one in $(w^A, \lambda^{-1/\sigma})$. (This is the homogeneity assumption referred to in sections 1 and 2, that is important for the result that a tax on direct household use of energy has no effect on non-energy output.) In our numerical work we furthermore specify the value of 2 for σ. As is noted in Table 1, this value implies that the elasticity of consumption growth (specifically, growth in consumption of the aggregate A) between two periods, with respect to the real rate of return between those periods (also measured in terms of the composite good A), holding hours worked constant, is equal to .5. This value (which follows Rotemberg and Woodford (1993)) is within the range of values consistent with a variety of studies of the relation between intertemporal substitution in consumption and asset prices. (A value of 1 is common in the real business cycle literature.)

The only features of the Frisch demands that matter for the log-linear approximation to the equilibrium conditions are the elasticities of the functions with respect to their two arguments, again evaluated at the steady-state equilibrium. However, the homogeneity assumption stated above implies that all four elasticities are uniquely determined once we specify values for σ and any one of the elasticities. We choose to calibrate the model in terms of a specified value for ϵ_{Hw}, the elasticity of the Frisch labor supply function with respect to the consumption wage, because this particular elasticity (sometimes called the "intertemporal elasticity of labor supply") is both familiar and the subject of a large number of econometric studies. The value that we

[22] For demonstration of how the equilibrium conditions can conveniently be written in terms of the Frisch demand functions, see Rotemberg and Woodford (1993). The discussion below of the parameterization of the Frisch demand functions follows Rotemberg and Woodford (1992).

[23] This is equivalent to a homogeneity assumption on the function U. For further discussion of the class of functions U satisfying this condition, see King, Plosser and Rebelo (1988) or Rotemberg and Woodford (1992).

assume (again following Rotemberg and Woodford (1993)) is at the high end of the range of values obtained from panel data studies, though it is considerably smaller than the values most often assumed in the real business cycle literature (often 4 or more).

The remaining feature of household preferences to specify is the aggregator function $A(C, E^h)$. Again, because we assume that the function is homogeneous degree one, the only parameters for which numerical values are needed are the elasticities of A with respect to its two arguments, and the elasticity of substitution between the two arguments, again evaluated at the steady-state equilibrium consumption bundle. The elasticities with respect to the arguments are again implied by the share parameters specified above. For example, the elasticity of A with respect to C is given by

$$\frac{C}{C + p_E E^h} = \frac{s_C(1 - s_M)}{s_C(1 - s_M) + \frac{\theta^h}{1-\theta^k} s_E}$$

Thus it remains only to specify ϵ_{CE}. Our value is taken from the econometric study by Houthakker, Verleger and Sheehan (1974).

We finally describe the parameters that specify the product market structure. As noted in the text, all of the models that we consider amount to different specifications of the markup function $\mu(X/Y)$ in equation 6. The features of this function that matter for the log-linear approximation to the equilibrium conditions are its value μ in the steady-state equilibrium, and the elasticity of the function with respect to its argument X/Y, also evaluated at the steady-state value of that argument. In the case of the competitive model, we specify $\mu = 1$ and $\epsilon_\mu = 0$. In the case of the monopolistically competitive (or "constant markup") model, we specify $\mu = 1.2$ and $\epsilon_\mu = 0$. In the case of the oligopolistic (or "variable markup") model, we specify $\mu = 1.2$ and $\epsilon_\mu = .15$. As we discuss further in Rotemberg and Woodford (1992), the amount of market power assumed in the steady state in the case of the imperfectly competitive specifications (prices 20% in excess of marginal cost) is within the range of estimates obtained by a number of studies of U.S. industries. In that same paper we show that the implicit collusion model implies theoretical bounds upon the value of ϵ_μ, namely, that $0 < \epsilon_\mu < \mu - 1$. The value that we assume here satisfies the theoretical bound. These parameter values for the implicit collusion model also coincide with those that are shown in Rotemberg and Woodford (1993) to predict effects of oil price shocks that are similar to those observed during the period 1947-1980.

In the case of the oligopolistic model, it is also necessary to specify a value for the parameter α which appears in the definition of X provided in Rotemberg and Woodford (1992). This parameter indicates the expected rate of growth of a given oligopoly's share in total expenditure. We assume $\alpha = .9$, because, as is discussed in the Appendix of Rotemberg and Woodford (1992), this value is consistent with the existence of an equilibrium with imperfect collusion (a binding incentive compatibility constraint) in the case of oligopolies with no more than ten firms.

Finally, in the case of either of the imperfectly competitive models, we must specify the parameter ρ in equation 13. We set this arbitrarily at .2. This parameter does not seem to have an important qualitative effect on our results as long as it is relatively small (adjustment of number of firms is not too fast).

Table 1
Energy Use in the U.S. Economy

		Quantity	Price		Value (billions)
Coal					
Production	1000	short tons			
Exports	100	short tons			
For Electricity	800	short tons			
Other uses	100	short tons	30	/ton	3
Petroleum					
Production	2707	Million Barrels	16.54	/barrel	44.77
Imports	2223	Million Barrels	16.54	/barrel	36.77
For electricity	196	Million Barrels	18.65	/barrel	3.70
Petroleum and Coal processing					
Value added					36.70
Natural Gas					
Total Revenues					69.07
For Electricity	2.78	Trillion Cu.Ft.	2200	/Cu.Ft.	6.10
Other uses					62.93
Electric Utilities					
Total revenues					184.90
Total Purchases					365.41

Table 2

The Calibrated Parameters

Parameter	Values	Description
g	.03	Rate of technical progress (per year)
r	.06	Steady state real rate of return (per year)
δ	.073	Rate of depreciation of capital stock (per year)
s_C	.697	Share of private consumption in final uses
s_I	.186	Share of private investment in final uses
s_G	.117	Share of government purchases in final uses
s_E	.02	Share of energy costs in total costs
s_M	.5	Share of materials costs in total costs
s_H	.36	Share of labor costs in total costs
s_K	.12	Share of capital costs in total costs
s_Φ	0, .167	Share of fixed costs in total costs
θ^{dom}	.83	Share of domestically produced energy in total energy use
θ^h	.4	Share of direct household use in total energy use
η^g	.17	Share of hours hired by the government
ϵ_{VG}	.69	Elasticity of substitution between value added and intermediate inputs
ϵ_{EM}	.18	Elasticity of substitution between energy and materials
ϵ_{KH}	1	Elasticity of substitution between capital and hours
$1/\sigma$	.5	Elasticity of intertemporal substitution of household expenditure
ϵ_{Hw}	1.3	Intertemporal elasticity of labor supply
μ	1, 1.2	Steady state markup (ratio of price to marginal cost)
ϵ_μ	0, .15	Elasticity of the markup with respect to X/Y
α	.9	Expected rate of growth of individual oligopoly's expenditure share
ρ	.2	Rate of partial adjustment of number of firms

Table 3

Long-Run Effects of an Energy Tax

	[TAX ON ALL ENERGY USE]		[HOUSEHOLD USE ONLY]	
	Competitive	Market Power	Competitive	Market Power
Non-Energy Output	-.071	-.097	0	0
GDP	-.072	-.098	-.022	-.022
Hours Worked	-.024	-.021	0	0
Product Wage	-.056	-.085	0	0
Consumption Wage	-.092	-.122	-.037	-.037
Energy Use in Prod.	-.271	-.297	0	0
Household Energy	-1.052	-1.082	-1.000	-1.000
Capital Stock	-.084	-.110	0	0
Number of Firms	—	-.033	—	0

Figure 1

Estimated Production Functions

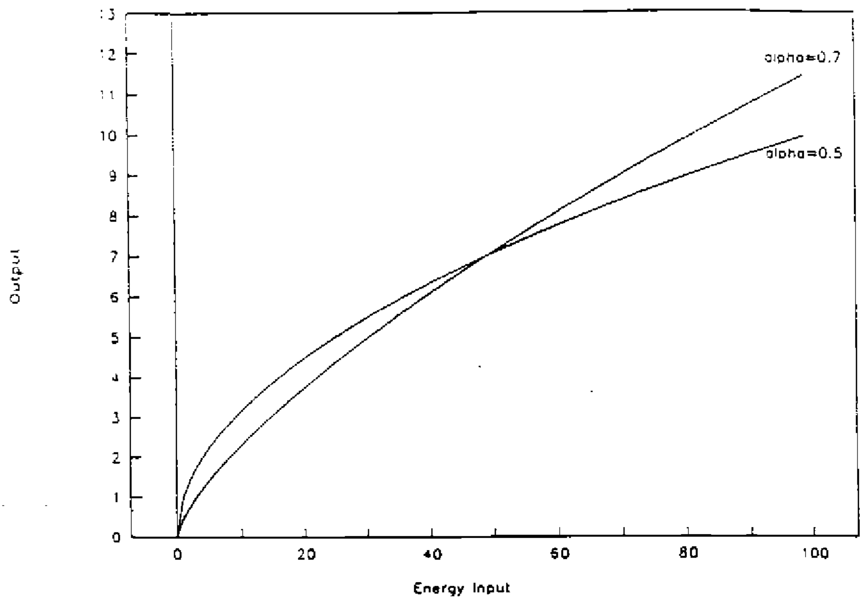

Figure 2

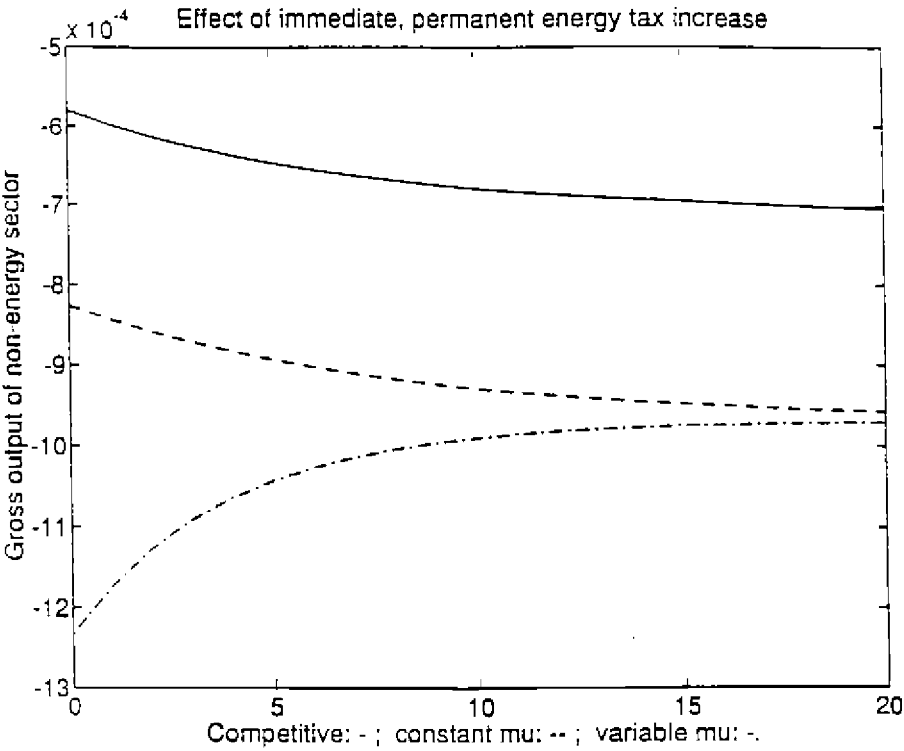
Effect of immediate, permanent energy tax increase
x 10^-4
-5
-6
-7
-8
-9
-10
-11
-12
-13
Gross output of non-energy sector
0
5
10
15
20
Competitive: - ; constant mu: -- ; variable mu: -.

Figure 3

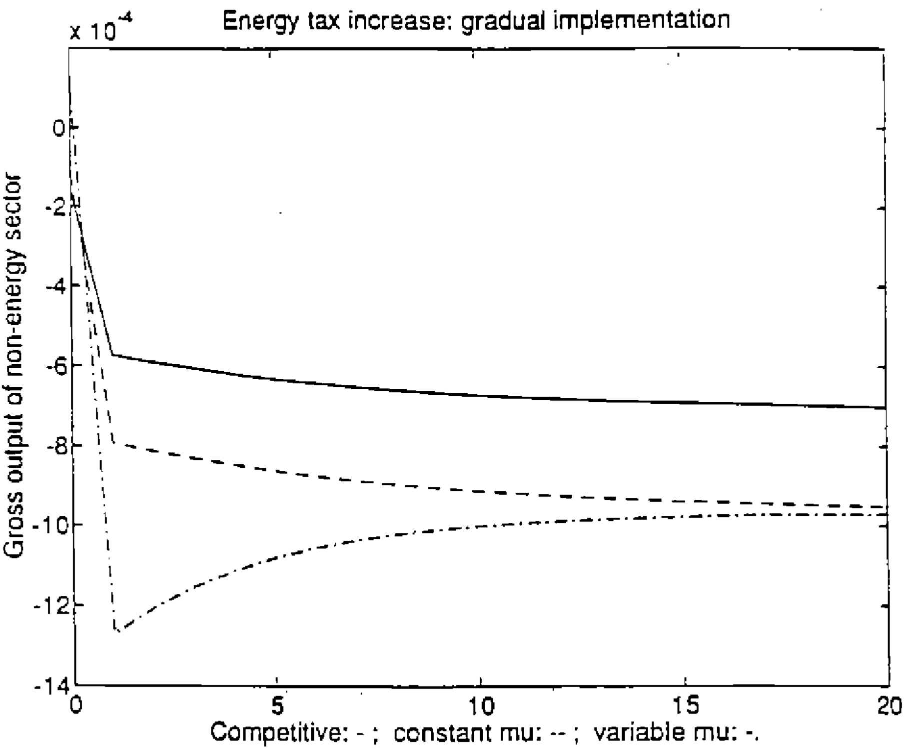

Energy tax increase: gradual implementation
x 10^-4
Gross output of non-energy sector
0
-2
-4
-6
-8
-10
-12
-14
0
5
10
15
20
Competitive: - ; constant mu: -- ; variable mu: -.

Figure 4

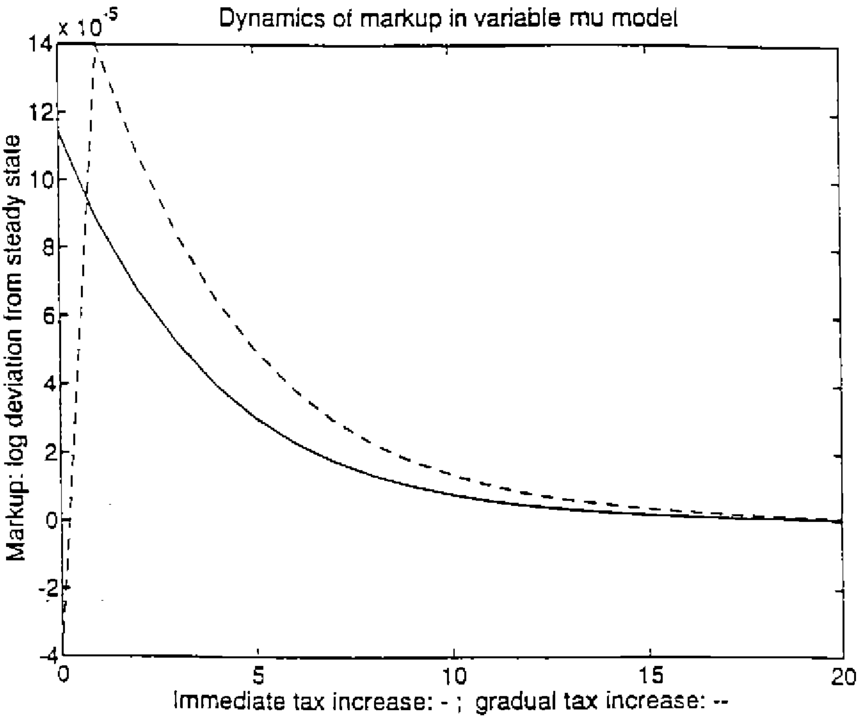

Dynamics of markup in variable mu model
x 10^-5
Markup: log deviation from steady state
14
12
10
8
6
4
2
0
-2
-4
0
5
10
15
20
Immediate tax increase: - ; gradual tax increase: --

Figure 5

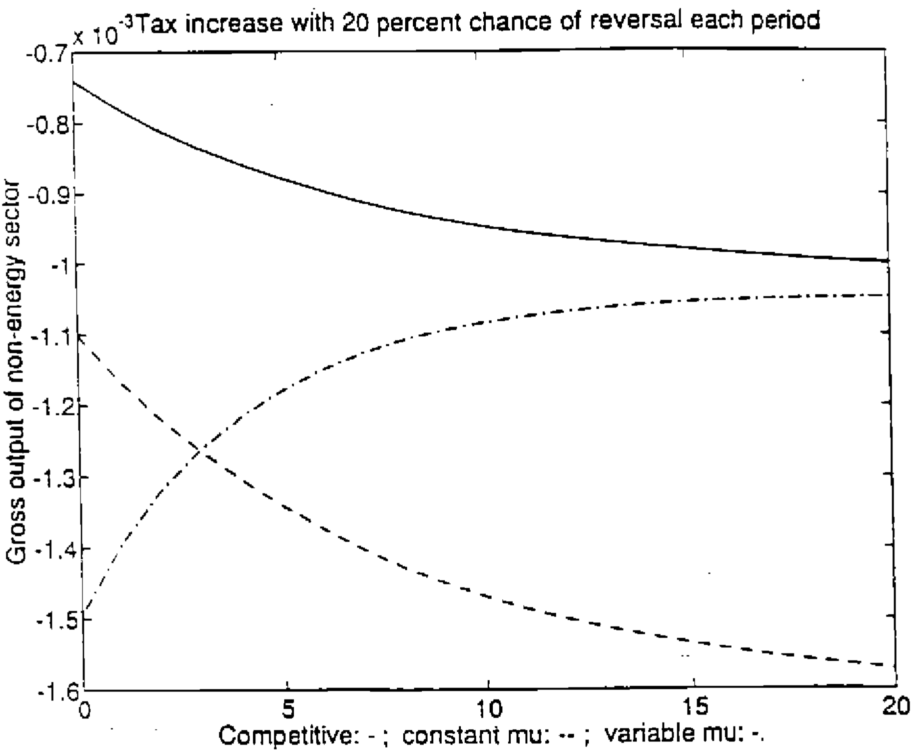

x 10^-3 Tax increase with 20 percent chance of reversal each period
-0.7
-0.8
-0.9
-1
-1.1
-1.2
-1.3
-1.4
-1.5
-1.6
Gross output of non-energy sector
0
5
10
15
20
Competitive: - ; constant mu: -- ; variable mu: -.