Modeling the Impact of an Energy Price Shock on Interregional Income Transfer

Clarence W. Nelson
Vice President and Economic Adviser
Research Department
Federal Reserve Bank of Minneapolis

In 1979–80 the Organization of Petroleum Exporting Countries (OPEC) doubled the world price of crude oil—the second upward oil price shock in a decade—and domestic producers of crude energy soon followed suit. Despite efforts to conserve energy, U.S. consumers' total energy spending increased as a result of the shock. Because certain parts of the United States have abundant energy resources while most other parts have little, consumers in energy-poor states felt that their energy-rich compatriots were reaping windfall gains—and largely at their expense.

Public representatives of the energy-importing states (those states consuming more energy than they produce) reinforced this perception.[1] In some discussions, analysts cited the energy-importing states' increase in energy spending as a measure of income siphoned away to crude energy producers—some to OPEC and much to the energy-exporting states.[2]

That measure of income drainage was sizable. For example, data on U.S. energy consumption and production in 1976 show that the energy-importing states (39 in all) bought about 7 billion barrel-equivalents of crude energy from other domestic suppliers or OPEC.[3] The 1979–80 price increase of roughly $8 per barrel-equivalent of energy would therefore have increased their spending on "imported" energy by $56 billion. Even if energy consumers responded to the price hike by reducing their demand by, say, 20 percent, the increase in their energy spending would still have been an impressive $45 billion.

That sizable measure of income transfer from energy-importing states to crude energy producers is incomplete, however, because it does not account for a number of factors in the U.S. economy that redistribute windfall income gains (or losses) across sectors and regions.[4] These factors include

- *The federal tax system.* Uncle Sam, through corporate and personal income taxes, takes a cut of windfall gains and then redistributes the revenue more evenly across the population.
- *Cross-regional ownership of stocks.* A Minnesotan holding Exxon stock receives higher dividends as the company prospers.
- *Cross-regional ownership of mineral rights.* A New Yorker who owns mineral rights to oil from a Louisiana well receives higher royalty payments.
- *Interregional trade.* Households in Texas have more

[1] In the popular press, accounts of debates over energy-related issues such as price decontrol and state severance taxes even alluded to the Civil War. See, for example, Wars between the states, 1981, p. 19, and Now energy is what counts in the war between the states, 1981, p. 166.

[2] See, for example, references to income transfer in Chase Econometrics 1978, p. C-6, and MDEPD 1982, p. 4.

[3] The 1976 U.S. energy consumption and production figures are based on data reported in Chase Econometrics 1978, Table 5, p. B-11.

[4] These redistributive factors were recognized earlier by Nelson (1981) and Pfister (1982), for example, as having a bearing on interstate income transfer resulting from an upward energy price shock.

disposable income to spend on cars produced in Michigan, thereby boosting Michigan household income a bit.

To measure the impact of an energy price shock on the interregional transfer of income, these factors have to be considered. That is what this paper attempts to do, using the 1979–80 energy price shock as a case study.

To begin, the paper provides a preliminary estimate of how much redistributive factors initially reduced income transferred from energy-importing to energy-exporting regions of the domestic economy as a result of the 1979–80 shock. This exercise introduces some key concepts used in the analysis and gives a sense of the magnitudes involved. The preliminary estimate, however, does not indicate how individual sectors of the economy would respond to the price shock, nor does it account for the economy's response over a longer stretch of time.

To address these additional issues, a formal, multi-sector model of the U.S. economy is presented. The model is set up to represent the relationships of product and payment flows among economic sectors in 1978, just before the shock. These relationships are then adjusted to reflect how the economy would be likely to respond to the 1979–80 shock over a couple years. A comparison of the model's pre- and postshock results shows that although the amount of income transferred from energy-importing to energy-exporting regions was significant, the redistributive factors did lessen the amount substantially. Moreover, the model's measure of income transfer is substantially lower than the increase in the energy-importing states' energy bill—an income loss measure often cited around the time of the shock.

A Preliminary Estimate of Income Transfer

This section presents a preliminary estimate of the amount by which the redistributive factors could lessen income transferred from energy-importing states to energy-exporting ones as a result of the 1979–80 shock. This exercise gives a sense of how these factors work to redistribute income, introduces the measure of income transfer used here, and provides a rough idea of the magnitudes involved.

Some Assumptions

To perform this exercise, some assumptions are needed to characterize the U.S. economy in 1978, just before the shock. I assume that the economy consists of two domestic regions, one with abundant crude energy supplies and the other with no crude energy resources at all. For convenience, I refer to the former as *energy-rich* and the latter as *energy-poor*. The two regions are an abstraction based on observations of the U.S. regional situation in 1978. Of the two, the energy-poor region is the larger, making up roughly five-sixths of the total U.S. population at the time.

Data on energy consumption and production show that total U.S. consumption in 1978 was the energy equivalent of about 13.8 billion barrels of petroleum.[5] Of this total, assume that 11.5 billion barrels are consumed by the energy-poor region and 2.3 billion by the energy-rich region. Given a 1978 refined energy price of $18.80 per barrel-equivalent, the energy-poor consumers spend $216.2 billion for refined energy and the energy-rich, $43.2 billion.

Of the 13.8 billion barrels consumed, 10.3 billion are assumed to be produced by the energy-rich region and 3.5 billion by foreign sources. I assume that the 1979–80 shock doubles the price of crude energy from $8 to $16 per barrel-equivalent, and as a result, the price of refined energy increases by $8, from $18.80 to $26.80. In other words, crude energy producers take the entire value increase per barrel, and energy refiners pass along the full per unit cost increase to consumers. This approximates what actually happened to energy prices after the 1979–80 shock.

U.S. consumers respond to the price hike by reducing their demand for energy by about 13 percent.[6] That puts their postshock energy consumption at 10 billion barrels for the energy-poor region and 2 billion barrels for the energy-rich. Multiplying these consumption figures by the higher price of refined energy ($26.80) gives a $62 billion increase in consumers' overall energy spending due to the shock—an energy bill nearly $52 billion higher for the energy-poor and over $10 billion higher for the energy-rich.

I assume that the change in energy demand is borne entirely by foreign producers so that domestic crude output remains the same as before the shock, at 10.3 billion barrels. This means that the domestic crude energy producers (located only in the energy-rich

---

[5]For the purposes of this paper, *energy* is not restricted to oil but is instead defined more broadly to include all forms—coal, natural gas, nuclear energy, and others. Energy is measured as an aggregate of its various physical forms, based on their total content of British thermal units (Btu). For convenience, all quantities of energy are converted to their equivalents in barrels of oil, using 5,880 thousand Btu per barrel as the conversion rate.

[6]The reduction in demand reflects an aggregate refined-energy demand elasticity of −0.3 percent, which says that for every one percentage point increase in the price of energy, demand for energy declines by three-tenths of a percent. This elasticity rate is consistent with some studies and also approximates what actually happened between 1979 and 1982.

region) enjoy a windfall revenue gain of $82.4 billion.

Now I assume that the entire windfall gain goes to households in the energy-rich region and that no redistributive factors operate in the economy. Given these assumptions, the relative effects of the price shock on the household income of each region are summarized as follows: The energy-rich region gains $82.4 billion from the production windfall while the energy-poor region gains nothing. After paying off the increase in their energy bills, the energy-rich households have a net gain of $72 billion left for nonenergy spending whereas the energy-poor households have a net loss of $51.8 billion.

A Measure of Income Transfer

Given these figures, how do I measure the amount of income transferred from the energy-poor region to the energy-rich region as a result of the price increase? The measure used here is the amount of money the energy-poor region would have had to transfer to the energy-rich region without the shock to cause the same difference in the two household sectors' income for nonenergy spending as with the shock. That is, the energy-rich households' income for nonenergy spending went up by $72 billion, while that of the energy-poor households went down by $51.8 billion, so the difference [$72 billion − ($−51.8 billion)] is $123.8 billion. This same difference would be accomplished without an energy shock if the energy-poor households handed over $61.9 billion to the energy-rich households, for then the energy-rich households' income would rise by $61.9 billion, while the energy-poor households' income would fall by the same amount. So the difference [$61.9 billion − ($−61.9 billion)] is also $123.8 billion. Thus, the measure used here is the difference in the two regions' nonenergy household spending, divided by two. This measure is intended to parallel the loss in consumer welfare in one region and the gain in the other following the energy price shock. Without redistributive factors operating, the value of this income transfer measure is a substantial $61.9 billion.

□ *With Federal Taxes*

I now calculate a measure of income transfer that includes the redistributive factor of federal income taxes. (The step-by-step calculations are shown in Table 1.) I again assume a windfall revenue gain of $82.4 billion to U.S. crude energy producers in the energy-rich region. From this windfall the producers deduct royalty payments and state severance taxes, which then accrue as income to the energy-rich region.[7] The producers' before-tax profits are then $65.5 billion. To these profits a federal corporate income tax is applied at the 1978 marginal rate of 52 percent, which gives $34.1 billion to Uncle Sam. The after-tax profits of $31.4 billion go as dividend payments to the energy-rich households, and these payments (plus royalties and severance tax gains totaling $16.9 billion) give them $48.3 billion. Now a personal income tax of 20 percent is applied to the income gain of the energy-rich households. The combined corporate and personal income tax revenues, totaling $43.8 billion, are redistributed among households in proportions roughly equivalent to population: one-sixth ($7.3 billion) to the energy-rich region and five-sixths ($36.5 billion) to the energy-poor region. From the resulting income, the increased energy spending is subtracted to give the amount of windfall revenue available for nonenergy spending.

From these results, the income transfer formula is applied: [$35.5 billion − ($−15.3 billion)] ÷ 2 = $25.4 billion. This amount, then, is the income transferred from the energy-poor region to the energy-rich one when the redistributive effects of federal corporate and personal income taxes are taken into account.

□ *With Taxes, Stock Dividends, and Royalties*

Calculations similar to those in Table 1 can be performed to incorporate the effects of cross-regional stock ownership and royalty holdings into the income transfer measure. In those calculations (not shown here) one-fourth of the royalty payments and two-thirds of the dividends accrue as income to the energy-poor region.[8] The resulting changes in nonenergy household spending are then used in the income transfer formula: [$16.8 billion − ($3.7 billion)] ÷ 2 = $6.6 billion.

This result shows that the cumulative effect of three of the redistributive factors is significant in reducing the amount of income transferred. Their significance is especially evident if this measure is compared with the original estimate of $61.9 billion when no redistributive factors were operant. The measure with the three factors reduces the original estimate by $55.3 billion.

More to Consider

The results of the preliminary estimate are useful in that

---

[7] Based on observations from 1978 data, I assume a flat royalty rate of 12.5 percent and a flat severance tax rate of 8 percent. Note that royalty payments and severance taxes do not work to redistribute income in the calculations of Table 1, but the two factors do affect taxes.

[8] The fractional shares are based on a reasonable assumption from observations on geographic patterns of royalty holdings and on a study of share ownership in sixteen U.S. oil companies (cited in Kalt and Leone 1984).

Table 1

Redistribution of Windfall Revenue Gains Via Federal Income Taxes

(Billions of Dollars)

Item	Crude Energy Producers' Accounts	Accrues as Income to Household Accounts in the	
		Energy-Rich Region	Energy-Poor Region
Windfall Revenue	$82.4	$0	$0
Tax Deductibles			
Royalty Payments (12.5%)	-10.3	10.3	0
Severance Taxes (8%)	-6.6	6.6	0
Before-Tax Profit	65.5	0	0
▶ Federal Corporate Income Tax (52%)	-34.1	0	0
After-Tax Profit	31.4	31.4	0
Change in Before-Tax Personal Income		48.3	0
▶ Personal Income Tax (20%)		-9.7	0
After-Tax Personal Income		38.6	0
▶ Transfer Payments from Federal Tax Revenue*		7.3	36.5
Gross Change in Personal Income		45.9	36.5
Increase in Postshock Energy Spending		-10.4	-51.8
Windfall Gain Available for Nonenergy Household Spending		35.5	-15.3

*Federal tax revenue totaling $43.8 billion ($34.1 billion plus $9.7 billion) are redistributed in proportions corresponding to population: one-sixth to the energy-rich region and five-sixths to the energy-poor region.

they provide a measure of the initial direct effects of the 1979–80 shock—a measure that considers three of the redistributive factors. But in that simple framework a number of important effects and interactions aren't handled adequately:

- *Interregional trade.* One issue needing consideration is the redistributive effect of interregional trade. Because the preliminary estimate is based on the assumption that all windfall gains are spent in the region in which they accrue, the estimate simply ignores trade. But people who receive windfall gains would be likely to spend some of their new income on goods produced outside their region. For instance, an oil magnate in Texas might use profits from oil to buy a car made in Detroit. So an estimate of income transfer should allow for changes in trade resulting from an energy price change.
- *Adjustments over time.* Another issue to consider is how people adjust to the price shock over time. The preliminary estimate provides a static accounting of the primary direct effects of the price shock. A fuller accounting of income transfer would need to con-

sider how people adjust their behavior in response to the shock and how these adjustments, over time, work their way through the economy. Among these adjustments are *substitution effects*—the adjustments producers and consumers make to a price increase by using other resources instead of the higher-priced commodity. For example, after an upward energy price shock, a trucking firm might tell its drivers to reduce truck speeds to 50 miles per hour to substitute more labor (driver hours) and capital (truck hours) for higher-priced diesel fuel. Or, a homeowner might substitute more capital (attic insulation) or goods (an extra sweater) for higher-priced heating fuel.

Despite their adjustments to conserve energy, consumers would still be likely to find themselves spending more dollars on energy and, as a result, would have that much less to spend on their region's business goods and services. Subsequently, lost sales would cause some worker layoffs and production cutbacks, which in turn would reduce consumer spending even further. This economic chain reaction multiplies the initial effects of a shock over time, and this ripple effect needs to be considered in measuring interregional income transfer.

- *Refinements.* Some additional refinements about differences in consumer behavior could also be considered. The preliminary estimate assumes that all consumers react to the price shock in the same way, reducing their energy demand by the same percentage. But 1978 data show that energy prices and the intensity of energy use vary among sectors and regions of the economy and that, accordingly, different sectors respond to a price increase by reducing their demand in different proportions. A fuller accounting could consider these differences.

To deal with these additional considerations, we need a way to systematically account for a number of economic sectors, the income and spending flows among them, and the adjustments to these flows that would be likely to occur as the result of a price shock. In short, we need a model.

The Model

This section presents a stylized model of the U.S. economy to analyze systematically the interregional income transfer resulting from an energy price shock.[9] The model is structured so that it can consider the four redistributive factors—federal taxes, stock ownership, royalty holdings, and interregional trade. In addition, the model also can account for the substitution effects, chain reactions, and differences in each sector's demand that could all be expected to accompany a price shock. I begin by describing each of the model's sectors; then I explain how the sectors are interrelated through a set of matrixes and prices and show how the model is solved.

Sectors

The model consists of thirteen sectors grouped into a foreign sector and the domestic economy, made up of a federal government sector and two regions—one energy-rich and the other energy-poor. (See the diagram of model sectors in Chart 1.) The two domestic regions each contain five sectors, the same five types in each. The energy-rich region, however, contains an additional sector, which is the only domestic sector that produces any crude energy.

The premise that all domestic crude energy production is confined to the energy-rich region means, of course, that neither the energy-rich nor the energy-poor region corresponds exactly to any actual grouping of states. Instead, the two domestic regions are stylized abstractions of the U.S. regional situation in 1978, just before the 1979–80 energy price shock.

To model the energy-rich region, I use state data for three key energy-producing states—Texas, Oklahoma, and Louisiana—as a template for determining the region's patterns of energy use. This seems reasonable, since those states produced about half of the nation's crude energy in 1978. As modeled, the energy-rich region can be loosely thought of as a composite U.S. region containing those three states as well as some other states and portions of states where crude energy production predominates. The region represents about one-sixth of the nation's population.

At a very basic level, the model's sectors are linked by flows of physical energy or other materials. Goods and services produced in one sector are purchased and used by other sectors and thus serve as intersector links. For instance, Chart 1 shows how sectors are connected by flows of crude and refined energy. For ease of exposition, the model's sectors can be grouped according to the energy flows as follows: producers of crude energy, refiners of crude energy, consumers of refined energy, and governments.[10]

---

[9]The model presented here draws upon earlier studies by Hudson and Jorgenson (1976), Krohm (1981), Nelson (1984), and Kalt and Leone (1984, 1986). A technical paper describing the model in detail will be available on request to the Research Department.

[10]A standard way of grouping sectors in input-output models like the one presented here is by *processing* and *final demand* sectors, as in Hudson and Jorgenson 1976. In my exposition, the processing sectors correspond to sectors

Chart 1 The Model's Sectors and Intersector Energy Flows

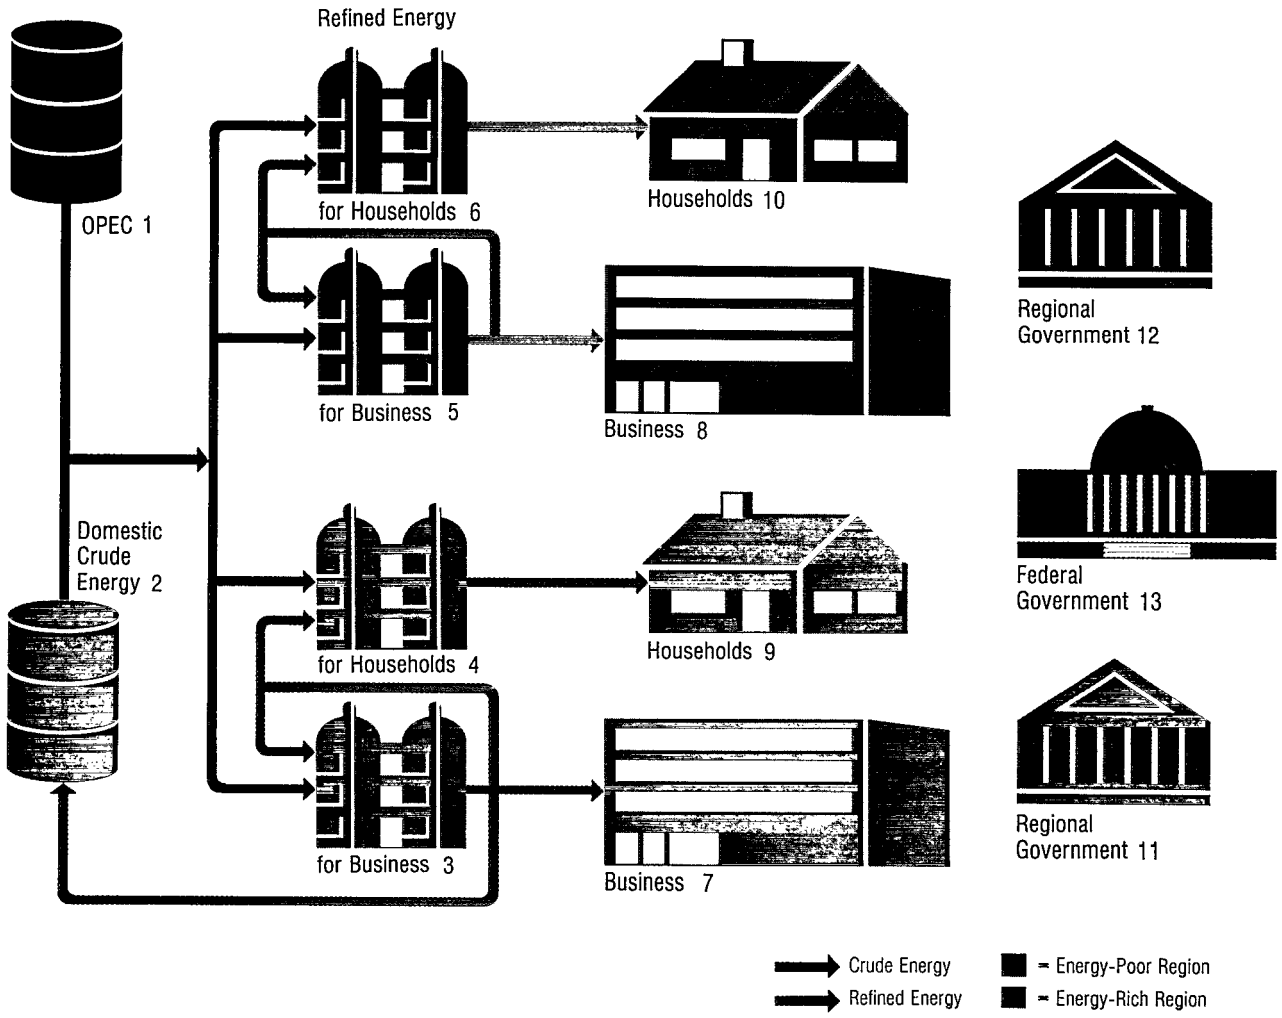

□ *Producers*

In the model, the U.S. economy's needs for crude energy are met by two energy-producing sectors: a foreign sector, simply called *OPEC,* and a domestic crude energy sector, located in the energy-rich region. Based on observations from 1978 data, I assume that these two sectors provided the U.S. economy with a total of 13.8 billion barrel-equivalents of energy. Of that amount, about 25 percent is assigned to OPEC and the remaining 75 percent to the domestic crude energy sector.

It is assumed that all of OPEC's proceeds from energy exported to the United States are used to buy goods and services from U.S. businesses. This trade is the only foreign trade explicitly modeled.[11] The domes-

3–8 and the final demand sectors to sectors 1,2, 9–13. The Hudson–Jorgenson grouping is used in performing the model's calculations.

[11]This assumption does not require that OPEC, of itself, spend its energy proceeds on U.S. business output; rather, it requires that OPEC plus other foreign buyers as a composite do so. Thus, trade is balanced, and the only foreign trade explicit in the model is trade financed by U.S. import dollars. Other foreign trade is assumed to be in balance and thus to net zero for U.S. output and consumption.

tic crude energy sector also buys goods and services from businesses in both the energy-rich and the energy-poor regions; in addition, it buys some refined energy from its own region's refiners.

☐ *Refiners*

Crude energy from OPEC and the domestic crude energy sector is processed into refined energy by four domestic refining sectors, two in each region. (For simplicity, it is assumed that one barrel of crude energy is required to produce one barrel of refined energy.) The refining sectors are divided into two types: those producing refined energy for business and those producing it for households. This split is an attempt to portray four different aggregate refining technologies to accommodate the observations, based on 1978 data, that for each energy unit, households pay more than businesses and energy-poor states pay more than energy-rich ones. The price disparity among sectors is explained by the value added in processing, packaging, transporting, and distributing refined energy as well as by the fact that energy-poor states tend to be located farther from crude energy sources than energy-rich states. Since different average prices would reflect different production costs, this difference needs to be accounted for in the model by the four separate refining technologies.

In their own production processes, the four refiners use some refined energy for processing. (Refined energy is also used for processing by the energy-rich region's crude energy sector.) Most of the refined energy, however, goes to the model's main energy consumers.

☐ *Consumers*

The four principal consumers of refined energy consist of a business and a household sector for each region. To model these four sectors, I regrouped data on 1978 energy spending and consumption, published by the Energy Information Administration (1978), to represent these two types of energy consumers. This division of consumers into businesses and households is necessary because changes in regional household spending are used to calculate the income transfer measure.

Business. The two business sectors produce all the nonenergy goods in the economy. In the production process, each uses refined energy from its region, labor from its region's household sector, and some business sector output (in the form of semiprocessed or intermediate goods) from itself and the other region's business sector. The goods produced by the two business sectors can be thought of as the primary stuff of gross national product (GNP).

Except for the federal government, all of the model's sectors buy products from the business sectors; the household sectors, however, patronize only the business sector in their own region. (In other words, when Texans buy cars made in Detroit, they make their purchases at a Texas dealership.) Because each business sector's output is exported to all but two sectors, the model can account for the interregional effects of an energy price shock as they feed back and forth between regions through changes in the level of trade.

As modeled, the two business sectors use a substantial part (roughly 50 percent) of their own output as production inputs.[12] This intrasector feature allows the model to account for the chain reaction effects on secondary or supporting industries when demand for primary business output changes.

Households. The household sectors, one per region, consume refined energy and business sector output. Households also play the role of supplying labor to the business and energy sectors within their own region, and for their labor they receive wages.

Besides earning income from wages, households in both regions receive dividend income by owning stock. As modeled, stock ownership in general corporations is distributed evenly among households in both regions. In contrast, households in the energy-rich region own twice as much stock in crude energy–producing companies, on average, as do households in the energy-poor region. These assumptions on stock ownership patterns, based on data from the New York Stock Exchange (1982) and shares used by Kalt and Leone (1984), are how the model accounts for the redistributive effect of cross-regional stock ownership.

Households in both regions also collect royalty payments from holding mineral rights to crude energy extracted in the energy-rich region. It is assumed, however, that energy-rich households receive 50 percent of all royalty payments made, whereas energy-poor households receive only 25 percent. (The remaining 25 percent goes to government.) This assumption seems reasonable, since it is likely that royalty holdings would be closely tied to early patterns of land ownership. This feature of the model thus takes into account the role of royalty holdings in redistributing income among regions.

---

[12]The ratio of intrasector-to-total output is based on generalized data from input-output models of the U.S. economy. See, for example, Yan 1969, p. 57.

☐ *Governments*

To account for the redistributive effects of taxes on interregional income transfer, the model includes three government sectors: two regional governments and a federal government. The revenues of the three sectors are provided primarily by taxes. Each regional government taxes all sectors within its region, and the federal government taxes all domestic sectors except regional governments. The federal government and the energy-rich region's government also collect some royalty payments from holding mineral rights to domestic crude energy. Based on information published at the time on the amount of royalties received by federal and state governments, I allocate 15 percent of all royalty payments to the federal government and 10 percent to the energy-rich region's government.

The spending behavior of the regional and federal governments differs. The regional governments spend some of their revenues for business sector output, mostly from their own region but some from the other one. The remaining revenues are distributed as direct transfer payments back to their own region's household sector. The federal government, in contrast, buys neither goods nor energy. Instead, it simply divides its revenue between the two regional governments according to some specified proportion and returns the revenue as intergovernmental transfer payments. In the model, the specified proportion is based on the population of each domestic region: one-sixth to the energy-rich and five-sixths to the energy-poor.

In reality, of course, governments use appreciable amounts of labor, energy, and other inputs to produce a portion of measured GNP. In the model, however, government energy and labor inputs are accounted for indirectly in the goods and services purchased from the two business sectors. Thus, as modeled, the government sectors are quite stylized abstractions intended, together, to reflect the income transfer aspect of government policies.

Relationships Among Sectors

The model's thirteen sectors have now been identified and briefly described in relationship to the flows of crude and refined energy. If, however, an attempt were made to depict not just energy flows but all intersector flows of products, labor, and payments, the resulting flowchart would be quite complex. A simpler, more convenient way of representing the model's structure, with all its intersector connections, is with matrixes, such as the one shown in Chart 2.

The matrix of Chart 2, consisting of rows and columns for each of the thirteen sectors, may be thought of as a general summary of all the model's intersector transactions. Each letter in a cell represents a transaction between two sectors (identified by row and column heads). In some cells, more than one letter appears, indicating different types of intersector transactions. When a cell has the same row and column number (that is, a cell located on the matrix's diagonal), then an intrasector transaction is indicated. If a cell is blank, then no relevant transaction occurs between the two sectors.

The matrix can be read in two ways, depending on the type of flows being tracked—flows of physical quantities or flows of dollar payments. On one level, it can be read as an *input-output quantities matrix,* called *matrix Q,* summarizing the physical flows of energy, business output, and labor among sectors. The sectors of rows 1–10 are the *sellers* of output to the *buyers* of columns 1–12. When reading across the rows, matrix *Q* traces the output from each sector as it becomes input to another: Output from the crude energy sectors (rows 1–2) becomes input to the four energy-refining sectors (columns 3–6). Output from the energy refiners (rows 3–6) is used in processing (by columns 2–6) or becomes input to the energy consumers (columns 7–10). Output from the business sectors (rows 7–8) goes to all sectors except the federal government (column 13) and the other region's household sector (column 9 or 10). The household sectors (rows 9–10) provide labor (hours worked) to the producing sectors of their respective regions (columns 2–8). [Note that the dividends, royalties, transfer payments, and taxes of the household and government sectors (rows 9–13) are purely financial transactions and are therefore excluded from the product flows of matrix *Q.*] For matrix *Q,* all quantities in a row can be summed across to determine a sector's total output. The columns, however, cannot be summed because their cells represent a variety of physical inputs (energy, goods and services, hours worked). Even so, the entries in each column can be thought of as a *recipe* of inputs used by each producing sector.

On a second level, Chart 2 can be read as a *payment flows matrix,* called *matrix F,* accounting for each sector's spending to and income from another sector. So rather than tracking quantities, matrix *F* tracks dollar payments representing sales and purchases of products as well as payments representing purely financial transactions. Because the payment flows move in the opposite direction from the product flows of matrix *Q,* matrix *F* is read down the columns. For

Chart 2 **Matrix of Intersector Flows**

Output From/Income to			*Input to/Spending by*	Producers		Refiners				Consumers				Governments			Total Income
				OPEC	ER Crude Energy	ER for Business	ER for Households	EP for Business	EP for Households	ER Business	EP Business	ER Households	EP Households	ER Regional	EP Regional	Federal	
				1	2	3	4	5	6	7	8	9	10	11	12	13	
Producers	OPEC	1				c	c	c	c								y
	ER Crude Energy	2				c	c	c	c								y
Refiners	ER for Business	3		e		e	e			e							y
	ER for Households	4										e					y
	EP for Business	5						e	e		e						y
	EP for Households	6											e				y
Consumers	ER Business	7	b	b		b	b	b	b	b	b	b		b	b		y
	EP Business	8	b	b		b	b	b	b	b	b		b	b	b		y
	ER Households	9		l d r		ld	ld	d	d	ld	d			g			y
	EP Households	10		dr		d	d	ld	ld	d	ld				g		y
Governments	ER Regional	11		t r		t	t			t		t				g	y
	EP Regional	12						t	t		t		t			g	y
	Federal	13		t r		t	t	t	t	t	t	t	t				y
	Total Spending		s	s		s	s	s	s	s	s	s	s	s	s	s	

ER = Enery-Rich Region
EP = Energy-Poor Region

Product Flows or Production Payments

c = crude energy
e = refined energy
b = business goods
l = labor

Financial Payments

d = dividends
r = royalties
t = taxes
g = government transfer payments

instance, reading down column 2 shows that the domestic crude energy sector spends money to buy refined energy (from row 3), business sector output (from rows 7–8), and labor (from row 9). After making all these production payments, sector 2 uses its remaining revenue to make financial payments: dividends (to rows 9–10), royalty payments (to rows 9–11 and 13), and taxes (to rows 11 and 13). The sum of all entries in a column gives the total spending s for that column's sector, and the sum of all payments in a row gives the total income y for that row's sector. In the model, each sector's total spending is set equal to its total income: $s = y$. This stipulation is in the spirit of general equilibrium modeling, which requires all sectors of an economy to be in a sustainable balance and all markets to clear.

To structure the model so that an energy price shock can be depicted systematically, two other matrixes are needed. The two can be constructed from matrixes Q and F, respectively, by expressing their values in relative terms. The first is derived from matrix Q by dividing each cell in a column by that column sector's total output (or row total). This operation, once performed for all cells, produces a *matrix of input-output ratios,* or *coefficients,* defining the relative sizes of all inputs and outputs. I refer to this matrix of input-output coefficients as *matrix A*. Each coefficient of matrix A says that to produce one unit of output from a column's sector, the quantity of inputs specified by the coefficient is needed from a given row's sector. Matrix A thus provides a useful summary of all the model's production relationships—that is, the *technology* of the modeled economy.

The second of the two matrixes can be derived in a similar manner by calculating relative payment flows from matrix F. Each cell of matrix F is divided by its column total. This operation, once performed for all cells, produces the *matrix of spending coefficients,* in which the individual elements (coefficients) are the fractional shares of total spending from a column sector to a row sector. I refer to this matrix of spending coefficients as *matrix B*. In matrix B, the columns must add up to one, since each cell represents a fraction of total spending for the column in which it appears.

Solving the Model

The four matrixes just described are related to each other by a set of relative prices of inputs and outputs p used by the producing sectors. The prices of crude energy and labor inputs are *exogenous* (taken as given from outside the model) whereas the prices of produced goods are *endogenous* (determined within the model) by the use of cost-markup equations.

The relationships between the price set and the matrixes can be expressed with simple algebraic notation. Assume the input-output coefficients of matrix A depend on relative prices: $A(p)$. Then let the spending coefficients in matrix B be denoted by

$$(1) \qquad B(p) = pA(p) + W(p)$$

where

$pA(p)$ = spending shares for production payments

$W(p)$ = spending shares for financial transactions (dividends, royalties, tax payments, and government transfers), also dependent on relative prices.

Let y be the total income received by each sector and let s be each sector's total spending. The relationship between y and s can be expressed as

$$(2) \qquad y = B(p)s.$$

Since this relationship holds for any level of income, an additional constraint is needed to set the income level at some benchmark. In the preshock version of the model, a constant k is added, corresponding to the total U.S. wage bill (the labor income paid to the two household sectors). The constant k also depends on prices: $k(p)$. Thus, the *model* is

$$(3) \qquad y = B(p)s + k(p).$$

Using the equilibrium condition that $y = s$, a *solution* to the model for a given p is a y^* such that

$$(4) \qquad y^* = B(p)y^* + k(p).$$

From equation (4) the payment flows of matrix F can be calculated:

$$(5) \qquad F(p) = B(p)y^*.$$

And from equation (5) the quantity flows of matrix Q can then be calculated:

$$(6) \qquad Q(p) = F(p)/p.$$

In short, this notation shows that if values can be as-

signed to the coefficients in matrixes A and B and to k, then values for sector incomes y and matrixes F and Q can be calculated, given a vector of prices p.[13]

Modeling the 1979–80 Energy Price Shock

Having explained how the model works in general terms, I now apply it to the 1979–80 energy price shock. Three steps are involved: First, the model's parameters and coefficients are set to reflect the preshock conditions of the U.S. economy, and the model is solved. Second, some model parameters and coefficients are changed to reflect how producers and consumers would adjust to the 1979–80 shock given a couple years' time; then the model is solved again. Third, the model's pre- and postshock solutions are compared; specifically, I compare the changes in the two household sectors' income available for nonenergy spending. From these figures, the measure of income transfer from energy-poor to energy-rich households can be calculated. In addition, the model can be used to perform some experiments to quantify how much each of the four redistributive factors contributed to the income transfer measure.

Before the Shock

To begin, the model needs to mimic, within its stylized structure, conditions in the U.S. economy just before the shock. To do this, data from 1978 or estimates from other studies are used to set parameters and specify coefficients in matrixes A and B.[14] But because not all values can be specified in this way, some components of the two matrixes remain as free parameters. These are *calibrated* by solving

$$(7) \qquad y_t = B(p_t)y_t + k(p_t)$$

where time $t = 1978$. That is, the free parameters are assigned values to produce sector income results that approximate the actual data for 1978 income flows. (For example, average tax rates are set so as to generate aggregate government revenues for 1978.) Once the values for all coefficients, prices, and total income levels are known, it is a simple matter, using equations (5) and (6), to compute the preshock values for the product flows in matrix Q and the payment flows in matrix F.

After the Shock

The next step involves setting a postshock price for crude energy, which in turn changes the model's parameters and coefficients to reflect how sectors would be expected to adjust to the shock. Essentially, the model works through the following scenario:

1. *OPEC sets a new world price for crude energy that doubles the price to U.S. refiners, and domestic crude energy producers follow suit.* In the model, the exogenous price of crude energy is raised from \$8 to \$16—an increase corresponding to what happened to that price in 1979–80.
2. *Businesses raise their prices to reflect higher energy and production costs.* In the model, cost-markup equations are applied by the producing sectors (refined energy and business) to set the absolute prices of their outputs.
3. *Energy refiners reduce the amount of crude energy imported from OPEC.* The ratio of imported-to-domestic crude energy, a parameter that adjusts some coefficients in matrix A, is lowered from 0.25 to 0.16, corresponding to observed data for imported energy within a few years after the shock. The data suggest that U.S. refiners preferred domestic over foreign crude energy available at the same price.
4. *As the relative price of energy increases, households and businesses reduce their consumption of refined energy.* Price elasticities of demand (model param-

---

[13]The notation used here is purposely streamlined to convey the basics of the model in simplest terms. In the actual model, for example, the $pA(p)$ in equation (1) is computed as the indicated matrix product only for the six processing sectors (sectors 3–8). The computed matrix is then augmented by relevant spending share parameters for the other seven sectors to construct a full 13 x 13 matrix matching the dimensions of $B(p)$. For the purposes of solving the model, the dimensions of system (2) are augmented to include equations that fix the level of employment hours for the two regions, so system (3) is actually larger than system (2). In system (6), p represents a six-element vector of prices, whereas in system (1) p is a 6 x 6 matrix of relative prices P (that is, sector-by-sector ratios of input prices to output prices). Some technical niceties (essential for computation purposes) are smoothed over to keep the notation simple, including the point that the product $pA(p)$ in (1) is actually element-by-element multiplication of two matrixes P and $A(p)$ of equal size. And the matrix-vector product $B(p)s$ in (2) technically requires the transpose s' of the row vector s in order to be properly defined. The matrix-vector product $B(p)y$ in (5) is technically an element-by-element multiplication of matrix $B(p)$ by column vector y, by which every element in the first row in $B(p)$ is multiplied by the first element in y, and so on for the remaining rows, to construct a new matrix $F(p)$ the same size as $B(p)$. Similarly, the quotient $F(p)/p$ in (6) is technically an element-by-element division of matrix $F(p)$ by column vector p, by which all the elements in the first row in matrix $F(p)$ are divided by the first element in p, and so on, to create a new matrix $Q(p)$ equal in size to $F(p)$. Finally, $F(p)$ in (6) is a 6 x 13 portion (the six rows representing the processing sectors) of the 13 x 13 matrix $F(p)$ in (5).

[14]For example, to assign values for the producing sectors' spending shares $pA(p)$, I examined an assortment of published aggregate data from 1978 income statements for major energy corporations and for corporations in general. Industry revenues were divided into expense categories roughly corresponding to the model's input-supplying sectors: crude energy (for refining sectors only), refined energy, nonenergy business inputs, and labor. The remaining expenditures were categorized as a profit margin, consisting of tax payments and after-tax profits returned to stockholders.

eters specifying how much a sector's consumption changes per unit change in relative price) are used to decrease relative consumption. For households, the energy consumed per dollar of disposable income is the parameter reduced. For the producing sectors, the coefficients for refined energy inputs in matrix A are reduced. The elasticities used vary for each consuming sector and are based on calculations from observed changes in the average price and consumption of energy from two to three years after the shock.

5. *Producers use relatively less energy and relatively more labor and other inputs in their production processes.* For the producing sectors, the input-output coefficients for energy in matrix A are reduced (as in step 4) whereas coefficients for labor and business inputs are increased to reflect the substitution away from energy. (The quantitative amounts of these adjustments are taken from parameters in Hudson and Jorgenson 1976.)

6. *The producers' input adjustments result in a less efficient combination of resources under the existing technology, so overall business output drops.* The total output of the two business sectors in matrix Q is decreased by using an input parameter that specifies the percentage of efficiency loss as a function of the overall price level increase. (The theoretical and empirical bases for this adjustment are found in Tatom 1981 and Miller, Supel, and Turner 1980, respectively.) The reduced quantity of aggregate real output for the business sectors serves as the basis for the constant k of the postshock model.

7. *As income levels change, tax revenues collected by state and federal governments change in proportion to their effective marginal tax rates.* In matrix B, spending coefficients related to average tax rates in $W(p)$ are adjusted according to the new income flows and actual marginal tax rates in 1982.

Once these adjustments are incorporated into the model, the postshock values for p', $A(p')$, $W(p')$, and $k(p')$ can be computed. Then the postshock version of the model can be solved:[15]

$$(8) \quad y'^* = B(p')y'^* + k(p')$$

$$\equiv [p'A(p') + W(p')]y'^* + k(p').$$

Results

Once the pre- and postshock versions of the model are solved, the solutions y^* and y'^* can be compared to see what has changed. Since the income transfer measure focuses on the effects of the shock on income flows, the changes between the pre- and postshock versions of the payments flow matrix F, reported in Table 2, are examined. The table shows that the shock increases the household income available for nonenergy spending (that is, for purchases of business sector output) by \$12 billion for the energy-rich region but decreases it by \$15.6 billion for the energy-poor one. (In real terms, the resulting changes in each region's nonenergy household spending are an increase of 0.7 percent for the energy-rich region and a decline of 6.7 percent for the energy-poor.) Applying the income transfer measure to these results gives [\$12 billion − (\$−15.6 billion)] ÷ 2 = \$13.8 billion.

The model's \$13.8 billion measure of income transfer is roughly on the same order of magnitude as the earlier, preliminary estimate of \$6.6 billion. Certainly, both the preliminary and model estimates are far less than those cited by some observers at the time of the shock. The model's estimate builds in much more detail than the preliminary estimate, however, by considering all four redistributive factors, adjustments over time, and differences among sectors.

Some Experiments

The model can be used to conduct some experiments that quantify how much each redistributive factor lessened the income transferred from energy-poor to energy-rich households. The general method of each experiment involves adjusting certain coefficients, when solving the postshock model, to suppress one of the redistributive factors in the model's results. To do this, selected postshock flows responsible for redistributing income are reset to their preshock levels. Solving the adjusted postshock model then provides one way of estimating what the change in postshock income flows would have been if the redistributive factor had not been at work.

The simple mathematical notation can be used to summarize the method of the experiments as follows: Take elements of the spending coefficients for financial payments $W(p)$ and call them $w(p)$. Let the preshock version of these elements be the same as the postshock version:

---

[15] Prices p' and the matrix $A(p')$ must be solved simultaneously, since product prices depend on A. Let p_Q = input prices and p_A = product prices. We have $A = A(p_A, p_Q)$ and $p_A = f(A, p_Q)$, where $p \equiv (p_A, p_Q)$ and f reflects a cost-markup equation. Given p'_Q, we solve for p'_A and $A(p'_A, p'_Q)$ using a simple iterative process.

(9) $w(p) = w(p')$.

Then let

$$(10)\quad W'(p') = \begin{cases} W(p') \text{ for elements not in } w(p) \\ W(p) \text{ for elements in } w(p). \end{cases}$$

The experiment is to compute an adjusted postshock solution

$$(11)\quad y''^{*} = [p'A(p') + W'(p')]y''^{*} + k(p').$$

The income transfer formula is then applied to the results of the y''^{*} solution to produce a measure that can be compared with the y'^{*} income transfer measure of $13.8 billion, calculated earlier.

The results of these experiments are reported in Chart 3. Since the redistributive factors interact, it is not appropriate to expect the results to add up as components of a total. The calculations do, however, provide some idea of the relative contribution of each factor in lessening the amount of income transferred between regions.

□ *Federal Taxes*

The first experiment attempts to quantify how much federal income taxes, both corporate and personal, lessened the income transfer measure from what it would have been without this factor. To do this, I adjust

Table 2

Changes in Payment Flows of Matrix *F* as a Result of the 1979–80 Energy Price Shock

(Billions of Dollars)

Spending by / *Income to*	Producers		Refiners				Consumers				Governments		
	OPEC (1)	ER Crude Energy (2)	ER for Business (3)	ER for Households (4)	EP for Business (5)	EP for Households (6)	ER Business (7)	EP Business (8)	ER Households (9)	EP Households (10)	ER Regional (11)	EP Regional (12)	Federal (13)
Producers													
(1) OPEC			0.4	0.4	0.3	1.6							
(2) ER Crude Energy			10.1	4.6	37.5	23.9							
Refiners													
(3) ER for Business		3.5	0.5	0.4			7.0						
(4) ER for Households									6.8				
(5) EP for Business					1.7	2.7		31.8					
(6) EP for Households										33.1			
Consumers													
(7) ER Business	0.4	0.4	-0.1	0.1	-0.1	0.0	6.4	1.6	12.0		5.8	0.2	
(8) EP Business	2.3	0.2	-0.1	0.1	-3.6	0.2	6.3	50.2		-15.6	3.8	22.2	
(9) ER Households		17.2	-0.5	0.2	0.4	0.4	3.7	0.1			5.2		
(10) EP Households		30.6	0.4	0.4	-3.0	1.3	1.3	-18.4				12.1	
Governments													
(11) ER Regional		5.2	0.1	0.1			0.4		3.2				5.8
(12) EP Regional					0.5	0.5		0.1		2.9			30.5
(13) Federal		19.0	0.6	0.5	2.5	2.5	1.6	0.6	4.7	4.3			
Total Spending	2.7	76.1	11.4	6.8	36.2	33.1	26.7	66.0	26.7	24.7	14.8	34.5	36.3

ER = Energy-Rich Region, EP = Energy-Poor Region

the postshock model by setting the federal marginal tax rates in matrix *B* to zero. Then I readjust the regional government tax rates so that their tax revenues equal what the total tax revenue for both regional and federal governments had been before the shock. This procedure ensures that government, as a whole, continues to collect the same amount of taxes from the tax-paying sectors under the experiment; however, all tax revenues are spent by the regional government in which the revenues were raised. Thus, the redistributive effect of federal taxes is eliminated.

With these adjustments in place, the model produces an income transfer measure of $37.5 billion. Subtracting the income transfer result of $13.8 billion with all factors operating gives $23.7 billion as the amount by which income transfer was reduced as a result of federal taxes.[16]

☐ *Stock Ownership*

The second experiment attempts to measure how much cross-regional stock ownership changed the amount of income transferred between regions. Recall that the model's initial parameters had stock ownership in general corporations to be widely dispersed geographically. In contrast, the stockholders of crude energy–producing companies were assumed to be likelier to reside in the energy-rich region, so twice as many of these shares were held by residents of the energy-rich region than by residents of the energy-poor one.

For the experiment, I reset the dividend-related coefficients in matrix *B* so that energy-poor households receive the same amount of dividends after the shock as they did before. Thus, any shock-induced changes in corporate profits (from sectors 2–8) are all rechanneled to stockholders in the energy-rich region.

Results of the experiment show that of the four redistributive factors, cross-regional stock ownership had the largest impact—$32.2 billion—on reducing the amount of income transferred from the energy-poor region to the energy-rich one.

☐ *Royalty Holdings*

The third experiment attempts to measure the amount by which royalty payments from cross-regional ownership of mineral rights lessen income transfer. For the experiment, I set the royalty payment shares between the two household sectors so that payments to energy-poor households remain the same after the shock as they were before. The results show that of the four factors, royalty holdings had the least impact on the income transfer measure, lessening it by only $2.1 billion.

Chart 3

Results of Model Experiments

Amount Each Redistributive Factor Lessened the Income Transferred from Energy-Poor to Energy-Rich Households

(Billions of Dollars)

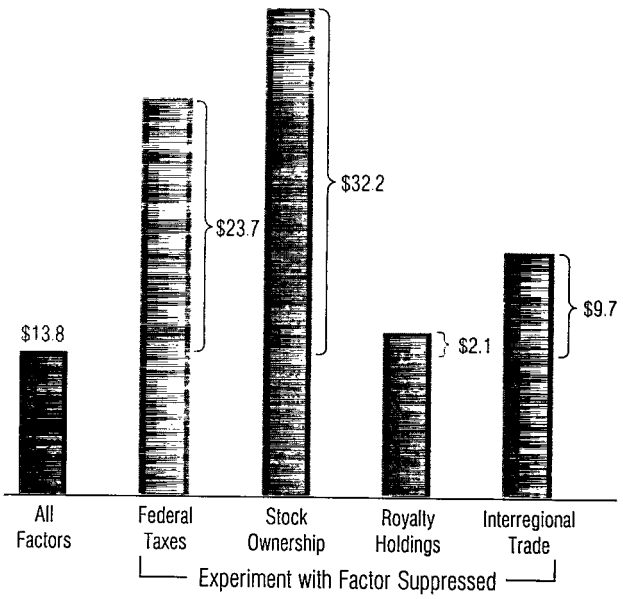

☐ *Interregional Trade*

The final experiment removes the interregional trade factor from the postshock version of the model to measure its impact on income transfer. In the model, all sectors except the federal government and each region's own household sector buy output from both business sectors. So any positive change in revenues from any producing sectors might be expected to spill over to the other region through expanded imports. (Conversely, any negative change in revenues might be expected to reduce imports from the other region.) The experiment eliminates the interregional trade factor by adjusting model parameters so that each sector's purchases from the other region's business sector remain

[16]The model experiments to measure the redistributive effects of the four factors do not attempt to incorporate the effects of the Crude Oil Windfall Profit Tax Act of 1980. I did, however, run a model simulation to include the effects of this tax. The results show an income transfer measure of $13.2 billion with the windfall tax included—an amount not appreciably different from the $13.8 billion result without it.

the same as under preshock conditions. Thus, changes in each sector's demand for inputs fall solely on its own region rather than acting directly to stimulate or depress business in the other region.

Results of this experiment show that the feedback effects of interregional trade served to reduce the income transfer measure by $9.7 billion—a much smaller reduction than those due to federal taxes or stock ownership. That this reduction is smaller, however, is hardly surprising, since interregional trade effects are generally a less direct path for rechanneling income flows.

Summary and Conclusions

This paper has shown how a stylized model of intersector product and payment flows can be used to analyze the interregional impact of an energy price shock. In analyzing the impact, the model takes into account the redistributive effects of the federal tax system, cross-regional ownership of stocks and mineral rights, and interregional trade. It also accounts for some substitution effects and chain reactions resulting from a shock. When applied to the 1979–80 upward energy price shock, the model's results show that the redistributive factors substantially lessened the amount of household income transferred from the energy-poor region to the energy-rich one.

Still, the nearly $14 billion in income transferred from the energy-poor households to the energy-rich ones is significant, for that measure represents a continuing, annual transfer of income from the energy-poor region to the energy-rich one. Expressed in present-value terms (assuming an interest rate of 10 percent), the stream of income transferred is equivalent to a one-time transfer of $140 billion. Since values of assets reflect the value of the income they earn, their values would be expected to change significantly in the two regions. To take a current example, the 1986 energy price crash has resulted in falling home prices in Texas and in turn has caused significant losses to that state's financial community.

Like most economic models, this one has its virtues and limitations. Among its virtues are the model's full accounting for the complex of intersector transactions on both spending and income sides; its systematic representation of sector technologies through the input-output matrix structure; its attempt to fuse the matrix modeling approach with a general equilibrium approach; and its structural detail, which is based on observed data.

Among the limitations of the calibrated model are the various stylizations and simplifications used to depict complicated economic relationships among sectors and regions. Examples of these limitations include the following:

- The two regions, as modeled, portray a worst-case scenario in energy resource disparity rather than corresponding to any actual groupings of states.
- The assumption that the OPEC sector spends all its increase in energy-dollar profits on U.S. business output is an oversimplification.
- The model generates an increase in the energy-rich region's labor hours and a decrease in the energy-poor region's hours (both ranging from 2 to 3 percent) but without accounting for the probable migration of labor between the two regions.
- Because the model uses an aggregate household sector for each region, it can't capture the differences in households' asset holdings and energy use that are observed across income groups.
- The model doesn't distinguish between investment and consumption, so it can't capture the investment booms or busts associated with an energy price shock.

These limitations suggest directions for improvements and future research. For instance, it might be desirable to divide the model's two household sectors into high- and low-income households, since upper-income households are likelier to hold stocks and are therefore likelier to benefit from the windfall gains of energy companies. Using data on the distribution of income and asset holdings, the results of this division might well show upper-income households in the energy-poor region enjoying real income gains and the lower-income households in the energy-rich region suffering income losses.

Another improvement to the model would be to incorporate business investment. As it stands, the model doesn't explicitly allow for the redistributive channel of investment to respond to a price change. However, investment to expand production capacity is an important component of production, and relative price changes would affect such investment. For instance, an upward energy price shock would be expected to stimulate investment in energy exploration and energy-economizing technology.

Although the model used here is calibrated to analyze a specific energy price shock, its framework could be extended and adjusted to other such shocks.

For instance, it would be relatively easy to adapt the model to analyze a downward energy price shock, such as the one in 1986. To do so, most of the processes modeled here could be reversed to produce an income transfer from the energy-rich to the energy-poor region. Of course, some model coefficients would need adjustment to account for changing patterns of economic behavior and the tax changes that have occurred since the 1978 conditions to which the model is calibrated. Even so, it is likely that the model's basic results would be similar, showing that the redistributive factors examined would substantially lessen the amount of income transferred between the two regions.

References

Chase Econometrics. 1978. State energy balances: Analysis and data. Material extracted from *Reports on the energy rich and the energy poor.*

Energy Information Administration. U.S. Department of Energy. 1978. *State energy data report.* Washington, D.C.

Hudson, Edward A., and Jorgenson, Dale W. 1976. Tax policy and energy conservation. In *Econometric Studies of U.S. Energy Policy,* ed. Dale W. Jorgenson, pp. 9–94. Amsterdam: North-Holland.

Kalt, Joseph P., and Leone, Robert A. 1984. A model of regional income accrual under energy price decontrol. Discussion Paper 1041. Harvard Institute for Economic Research.

_____. 1986. Regional effects of energy price decontrol: The roles of interregional trade, stockholding, and microeconomic incidence. *Rand Journal of Economics* 17 (Summer): 201–13.

Krohm, Gregory. 1981. The connection between rising energy costs and regional income growth. Wisconsin Division of State Energy. Photocopy.

Miller, Preston J.; Supel, Thomas M.; and Turner, Thomas H. 1980. Estimating the effects of the oil-price shock. *Federal Reserve Bank of Minneapolis Quarterly Review* 4 (Winter): 10–17.

Minnesota Department of Energy Planning and Development (MDEPD). Energy Division. 1982. Energy policy and conservation report.

Nelson, Clarence W. 1981. Regional impact of rising energy prices. Special Studies. Federal Reserve Bank of Minneapolis.

_____. 1984. On analyzing interregional transfers of wealth caused by energy price shocks. Research Department Working Paper 256. Federal Reserve Bank of Minneapolis.

New York Stock Exchange. 1982. *Fact Book 1982.* New York: New York Stock Exchange.

Now energy is what counts in the war between the states. 1981. *Business Week,* October 26, pp. 166–70.

Pfister, Richard L. 1982. Regional impacts of rising energy prices: A skeptical view of the dire predictions. Paper prepared for the Annual Meeting of the Midcontinent Section, Regional Science Association, June 17–19, at Vail, Colorado.

Tatom, John A. 1981. Energy prices and short-run economic performance. *Review* 63 (January): 3–17. Federal Reserve Bank of St. Louis.

Wars between the states. 1981. *Time,* August 24, p. 19.

Yan, Chiou-shuang. 1969. *Introduction to input-output economics.* New York: Holt, Rinehart and Winston.

The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.