Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Energy Modeling for Policy Studies

William W. Hogan,

ENERGY MODELING FOR POLICY STUDIES

WILLIAM W. HOGAN

John F. Kennedy School of Government, Harvard University, Cambridge, Massachusetts 02138, william_hogan@harvard.edu

Energy policy modeling owes a great debt to the disciplines of operations research. Valuable modeling tools were available when the energy crisis struck unexpectedly. In turn, the immediate response to problem-driven policy modeling produced methodological challenges and innovations that have application outside the domain of energy. The early days of the explosive growth of energy modeling for policy studies provide illustrations of the interaction of problem identification, model formulation, problem analysis, and policy implementation in the tradition of operations research.

Energy modeling for policy studies exploded in the United States after the Arab oil embargo of 1973. The standard tool kit of operations research, especially mathematical programming, played a prominent role in the early days and the subsequent analytical history over the rest of the last century. The field is too extensive for this memoir to be in any way comprehensive. The focus here is on early developments not well documented elsewhere and work that spawned a body of analysis that was still growing at the beginning of the new millenium. A recurring theme is the importance of energy modeling as process more than energy models as products.

For those interested, the decade following the oil embargo has been the subject of other investigations of the role of the formal analysis in the policy process (Greenberger 1983; for international comparisons, see Baumgartner and Middtun 1987). The analytical lessons derived from models and forecasts have been examined (Hogan 1996). Details of the important and cumulative investment in energy modeling within the U.S. Department of Energy continue to appear. (See, for example, Murphy et al. 1988, Gabriel et al. 2001.) The work of the Energy Modeling Forum (EMF), operating at Stanford University since 1978, provides an extensive record of comparative modeling studies across topics ranging from "Energy and the Economy" to "Fuel Diversity in Electricity Markets." (For a useful summary, see the EMF website, *www.Stanford.edu/group/EMF.*)

The period I address begins around the time of the oil crisis and extends through the few years that followed, with an emphasis on public policy in the United States developing and applying energy models.

1. OIL CRISIS

Energy policy in the early 1970s was fully entangled with the legacy of U.S. wage and price controls. Natural gas prices had been regulated for some time, and oil prices were included as part of the general wage and price controls of 1971. Not surprisingly, the controls produced unintended results in the form of growing concern with looming shortages. One of the first influential applications of energy modeling appeared in work at MIT that used econometric simulations to illustrate the impact of price controls on natural gas (MacAvoy and Pindyck 1973). Although the work was controversial in 1973, it did penetrate the discussions in the White House, which was more than usually distracted by the Watergate scandal. However, a full embrace of the natural gas debate was years in the future. Energy policy was an increasing area of concern, but the growing rumblings were more about imminent shortages of oil. Oil production had peaked in the United States in 1970. There was pressure to increase oil imports despite the existence of an oil import quota system.

Following a series of internal studies, a Special Energy Committee in the White House recommended a number of changes in both policy and organization. On April 18, 1973, President Nixon acted to terminate the oil import quotas and replace these with a set of modest fees, thereby validating the already apparent trend of growing oil imports. A new Energy Policy Office was created in the White House, to be the home of the first of a string of energy czars. (He was Governor John Love of Colorado, to be followed by Charles DiBona, William Simon, Frank Zarb, John Sawhill, and on down through the Federal Energy Administration and then the Department of Energy.) In addition, the Department of Interior "...was directed to develop a capacity for gathering and analyzing energy data...." (Greenberger 1983, p. 104). Eric Zausner was recruited to Interior from the White House Council on Environmental Quality to bring fresh human energy to the task.

The Interior Department was already the home of the Bureau of Mines (BOM) and the Office of Oil and Gas (OOG), which included many people with long experience in these industries. However, there was no tradition of analysis of the type deemed necessary, and Zausner soon created a new organization, the Office of Energy Data and Analysis (OEDA), that he charged with expanding and improving the nation's capability to address the vaguely defined but growing energy problems. In September of

Subject classification: Professional: comments on.
Area of review: Anniversary Issue (Special).

0030-364X/02/5001-0089 $05.00
1526-5463 electronic ISSN

Operations Research © 2002 INFORMS
Vol. 50, No. 1, January–February 2002, pp. 89–95

1973, I joined this group which then still numbered in single digits. (Bart Holaday was head of the group; David Wood and Frank Alessio were the other analysts.) We spent our first few days meeting with the established experts in OOG, who took due note of our lack of any industry background. Our mandate was vague, and there was no imminent crisis; hence, defense of their bureaucratic turf was the issue at hand in these early meetings.

Clarity arrived unexpectedly with the Yom Kippur War and the Arab oil embargo of October 1973. Suddenly oil supplies from the Persian Gulf were curtailed, world oil prices jumped to unprecedented levels, and energy policy mutated from an afterthought to dealing with a first-order security crisis. The White House, Congress, and the press wanted information, immediately. They looked in the phone book and found the number of the OEDA. We were about to learn of one definition of an expert as someone who was "two weeks ahead."

After the initial chaos of briefings and emergency meetings in the White House, Zausner asked me to untangle the story of oil import country sources to produce an analysis that would connect sources and uses of oil. In particular, the task was to resolve the conflicting estimates being reported in the press about the likely impact of the oil embargo. After much inquiry with reporters and others, it turned out that the different estimates had all originated with the same individual in the OOG. He was an expert who clearly knew more than I did about oil imports, but just as clearly he had not developed the habit of remembering, much less writing down, his train of thought.

The first energy-modeling initiative of the OEDA then became the application of the discipline of using tables with row and column totals. The OOG experts provided various data and estimates, and we wrote down the assumptions, gathered the details into tables, and checked the totals for consistency. The mathematical tool required was addition. The OOG experts forgot their turf concerns and were suddenly happy to do the core dump and leave to us the subsequent analysis. Turf issues evaporated in contemplation of responsibility for presentation to the voracious press. At a hurried meeting to decide who should handle the briefing, the leader of the group recently defending their turf announced: "Bill, I think you are ready!" The result was the first government report on the implications of the oil embargo and the forerunner of the subsequent analyses and the continuing *Monthly Energy Review* still published by the Department of Energy (U.S. Department of Interior 1973).

Complete reliance on the memory or judgment of the experts was both unsatisfactory and unworkable. Fortunately, we found that the BOM had extensive data on historical U.S. energy production and consumption, with regional detail and fuel specifics. Unfortunately, this database consisted of old paper reports accumulated over the years and piled in the corner of one office. In parallel with the effort to capture expert opinion, therefore, the OEDA launched a project to retrieve the database from the BOM. On an emergency sole-source basis, we hired a contractor to keypunch the data on IBM punch cards and load it into a new computer time-sharing system where we would be able to do simple analyses, estimations, and manipulation. Regrettably, since the data were in the public domain, we did not think to negotiate any special access to what grew into the Data Resources Incorporated (DRI) energy database. I shudder to think how many millions of dollars the U.S. government spent in subsequent years to use these data.

The DRI database and regression software became an immediate mainstay for analyzing likely demand and supply conditions in the coming months. The early work was primitive, but it did highlight the importance of prices and price impacts in analyzing the alternative outcomes. Although national policy still included domestic price controls, the early modeling work emphasized the interplay between shortages and prices. On the streets, citizens were angry about the immediate effects of gasoline allocations and the resulting lines at the gas stations on "odd and even" days, where eligibility to buy gasoline depended on the number on your automobile license plate. Meanwhile, back in their offices, the analysts in OEDA were worried about how to get more and better price data and how to incorporate the effects in the models.

2. PROJECT INDEPENDENCE

As the immediate chaos of the oil embargo subsided, the effort to deal with its implications expanded rapidly. New bureaucracies were created to deal with oil allocations and price controls. At the same time, Zausner took on responsibility both for creating a much expanded energy analytical capability and for charting the strategic direction for the U.S. government. The structure morphed through various emergency committees under the White House, to the Federal Energy Office, the Federal Energy Administration, and eventually joined the Energy Research and Development Administration (ERDA) to form the Department of Energy.

The immediate coordinating effort was the creation of Project Independence (PI), set out by President Nixon in a speech of November 7, 1973. Likened to the Manhattan Project, the stated goal was to achieve energy independence by 1980. Soon Zausner was given the task of leading the effort to define the details of policy to achieve this as yet unexamined objective.

Problems of natural resource limitation and energy insecurity were in the air (Meadows et al. 1972). There had been previous examinations of U.S. energy policy. (See, for example, National Petroleum Council 1972, Dupree and West 1972). However, the politics of the time were not conducive to relying on the work already done by the industry. The paranoia of the day saw industry conspiracies in the oil embargo, with rumors of oil tankers being held out of port in order to increase the shortage. The government might use some of the tools developed by the industry, but the government would have to rely on its own analysis.

As for the work by the government itself, it was simply inadequate for the task at hand. The government reports summarized the expertise of the authors, but provided no tools for analyzing alternative policies. Either the assumptions were outrun by the changed condition in the world oil market, or the hurried efforts were too conditioned by the vested policy position of the agencies (Ray 1973). Zausner was going to have to do something new, and he quickly used the clout of his office to assemble a series of task forces from across the government with a membership that soon included hundreds of full-time analysts.

The idea was that these task forces would be sent out to scour the country for information about the many components of the energy system. There were task forces for oil and for synthetic fuels, for energy demand and environment and conservation, for oil refining and for electric utilities. This large effort presented a major conceptual and practical problem. The previous approach to doing such studies relied very much on the work of a small number of experts that would pull together disparate sources of information and produce an answer that synthesized many assumptions and judgments. Few of the assumptions or judgments were explicit or formalized. Hence this approach would not lend itself to the kind of rapid "what if" policy analysis that Zausner envisioned.

Furthermore, the scale of the effort created its own methodological challenge. Before, the task of the expert was to integrate all the assumptions and judgments needed to produce the forecast of the outcome. How much oil would be produced in 1980? Now the responsibility of the oil task force was changed to describing the information about oil so that it could later be integrated with the analogous information about coal, and conservation, and everything else that would affect the production of oil. The outputs of the task forces were not to be the outputs of the study; they were to be the inputs to the something else.

It soon became clear that the "something else" would be an energy model. The previous experience with oil import analysis convinced everyone that somehow this would have to be a model that combined economic principles with engineering detail. In particular, when looking over any substantial period with markedly changed conditions, it would be critical to account for the economic effects of prices as well as economic growth. The latter had always been considered, but the former would be a novelty. Furthermore, it was clear that the energy model would need regional detail and many entry points for analyzing alternative policy proposals. It would be a large energy model, and there was nothing like it available at the time.

The growing PI team included experts with a certain familiarity with business and engineering applications of energy models. For example, the oil industry had long since embraced earlier work using linear programming tools to model oil refinery operations (Charnes et al. 1952, Manne 1958). This work was well understood, and there was commercial or industry software that could be used to handle the technology of refineries. By the same token, the refinery input-output framework could be adapted to other conversion technologies such as electric power plants.

In addition to the engineers and business experts, the PI team recruited a number of economists to handle the task of developing explicit models of fuel demands as a function of prices, demographic factors, and economic growth. Their natural approach, building on the early development of the DRI data, was to produce econometric models of fuel demands using multiple regression analysis. The national data provided the elasticities, and the models were disaggregated by census regions.

Separate groups analyzed each technology and the transportation links for moving energy products from the producing regions to the consuming regions. It took a bit of time for the task forces to appreciate that they were not to produce the answer, but the inputs to the answer. The typical form of the inputs would be in supply and demand curves, with accompanying technical transformation coefficients. To handle dynamics, the simplistic approach was to posit a trajectory for prices and then extract the snapshot for the respective year to obtain an interim static analysis that would be consistent with the smooth trajectory.

The industry oil refinery models had long been used to produce shadow prices for the outputs, and these shadow prices could be connected to the market prices for the products. Furthermore, there was a general familiarity with the basic ideas of market spatial equilibrium in a network system (Samuelson 1952, Takayama and Judge 1971). It was clear that for a given level of fuel demand, the supply curves could be combined with the network representation in a linear program and the least-cost solution would be consistent with competitive equilibrium. Furthermore, the shadow prices for the demand constraints could be interpreted as the prices on the supply curve.

The argument could be extended further to include the demand side of the problem if the demand curves produced by the econometric modelers satisfied the integrability condition of symmetric derivatives. Given integrable (inverse) demand and supply curves (p_D, p_S) for vectors (d, s) of regional energy products, costs of various transformation and transport activities (c), and the technology coefficients in the matrices A, B, and T, the market equilibrium would be consistent with the maximization of social surplus:

$$\max_{x,d,s} \int_0^d p_D(z)\,dz - \int_0^s p_S(z)\,dz - cx$$

s.t.

$$Ax = d,$$

$$Bx = s,$$

$$Tx \leqslant t.$$

Using a standard piecewise linear approximation to the various demand and supply curves at different points in the network, this optimization problem could all be reduced to a linear program. The basic framework is shown in Figure 1 taken from the PI report. This would be a problem of high dimensionality for the time, but it would be within the scope of both computers and software to solve the model and characterize the competitive equilibrium. Zausner would be able to change assumptions about costs and technology in the morning and review the results before lunch. Integration and analysis would be easy, once the task forces completed their work.

There were two obstacles to this happy scenario. First, the demand model did not satisfy the assumptions. The econometricians were estimating partial equilibrium demand models that were not naturally separable. Hence, there was no reason to assume that they were integrable. Therefore, at best this would be an equilibrium problem, but not an optimization problem. The supply model could still combine all the various fuels in:

$$\min_{x,s} \int_0^s p_S(z)\,dz + cx$$

s.t.

$$Ax = d,$$

$$Bx = s,$$

$$Tx \leqslant t.$$

With $\pi(d)$ as the shadow prices for the demand constraints, the resulting equilibrium solution would need to satisfy:

$$p_D(d) = \pi(d).$$

This would be more complicated than the straight optimization problem. The approach developed was an iterative method that picked a reference price vector, formulated separable (and therefore integrable) approximations to the demand curves, and solved the social welfare approximation to get a new estimate of the supply price vector. (There were prices for seven final demand products in nine regions for a price vector of dimension 63.) This was then repeated in the hope of obtaining convergence to the underlying equilibrium problem. In effect, this incorporated the own-price demand elasticities in the welfare maximization model, but ignored the less important cross-price effects.

Susan Shaw and I did a quick test, using manual communication between computers, to see if this method would work in practice and found that convergence was rapid. In effect, we could use linear programming as a subroutine in the search for an equilibrium solution that would account for the joint own- and cross-price effects. This became known as the Project Independence Evaluation System (PIES) algorithm (Hogan 1975). Later, the source of the rapid convergence was suggested by analogy to the well-known Jacobi method for solving nonlinear equations (Ahn and Hogan 1982). The success of this algorithm in providing rapid solution for a large practical problem reinforced interest in such modeling applications. "To a large extent, the PIES model and the associated PIES algorithm have provided impetus for the growth of the field of finite-dimensional variational inequality and nonlinear complementarity problems" (Harker and Pang 1990, p. 162).

The second obstacle was more mundane, but no less important for implementation of the analysis. This was

Figure 1. Project independence modeling framework.

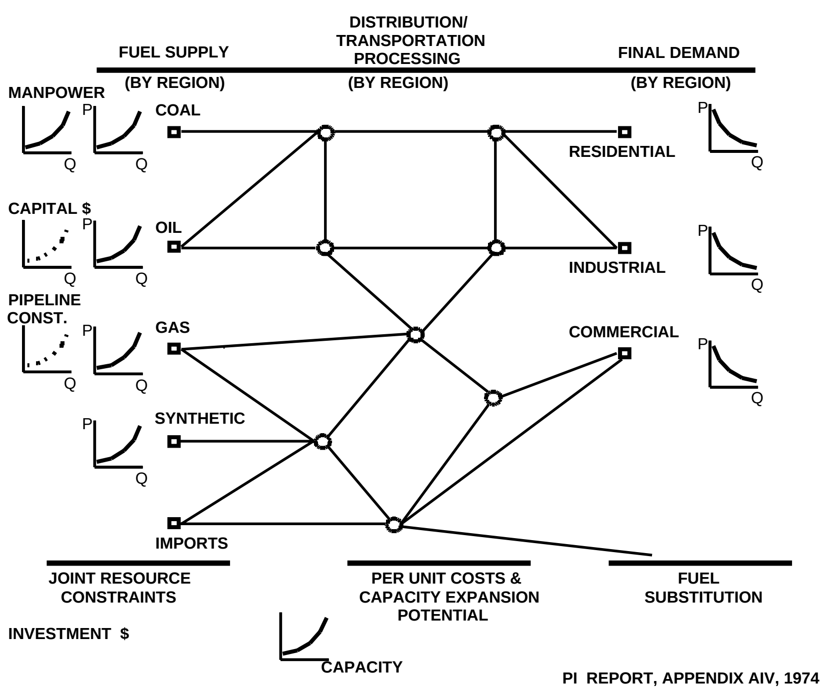

a time when computers were expensive and people were cheap. With a little tuning, the required computations were within the capability of the software and machines, but only barely as it turned out. Solving the linear program and applying the iterative PIES algorithm would be the easy part for the analysts. Getting the data into the standard form needed for the optimizer, and then reporting the results from the solution, would be the hard part.

This was before the days of sophisticated model generators like the Generalized Algebraic Modeling System (GAMS) that I use now. (For further information, see their website *www.gams.com*.) In 1974, we employed a version of a table-based matrix generator that had been developed for oil refinery models. But the PIES model required an extension of its matrix generation capabilities, and the model was large and complex. In retrospect, today we would say that we were debugging a beta version of the matrix generator. It took an enormous amount of scarce time and talent from dedicated public servants to get the data in and results out. (The Integration Task Force for the PIES model included J. D. Pearson, R. T. Eynon, M. H. Wagner, W. C. Mylander, Susan H. Shaw, now Susan Holte, and M. G. Rackoff.) It was agony, and subsequently I wrote a paper arguing that large models were harder to generate than to solve.

Years later I was back at the U.S. Department of Energy and reading some of their computer output from the descendants of PIES. The report was labeled "Wonder Bread." Upon inquiry, I found that the origin of the name had been lost in the oral history. I recounted our frustrations with the early report-writing software. When the original PIES team got a small part of the summary tables working, I breathed relief with the comment that "half a loaf is better than none." When the full report writer finally worked, much later, it was duly named after a then-popular brand of bread.

The many uses of the PIES model over the next few months became as intended. The changed input would be made in the morning in response to the policy interest, and the runs would be back by lunch. In the course of doing many debugging runs and policy scenarios, it seemed that we were using the combined processing power of several Control Data machines scattered across the nation. Mayor Daley once obliged by delaying street repairs in Chicago so that the circuits would not be interrupted. As I recall, there was a million dollar overrun on a two hundred thousand dollar computer time contract.

An example of the use of the model was in the investigation of the role of environmental constraints. The United States has much coal, and the supply curves reflected that costs rose modestly with increased use. In an effort to limit coal use, we tested a more stringent application of power plant emission standards. The model runs reported that coal consumption increased! Everyone was shocked, but the analysis revealed the explanation. The environmental restrictions were in the form of new source performance standards. When these were tightened, the economics tilted towards expanding the use and extending the life of old coal plants, which were inefficient and used more coal. After the fact, the answer was obvious, a clear indication of a good model and a good modeling process. The insight was the output of importance, not the numbers.

As it happened, the PI report organized a vast amount of information that showed that independence was a goal not to be obtained any time soon, and certainly not by 1980 (Federal Energy Administration 1974). Energy independence was not like a Manhattan Project under control of the government; the economics overwhelmed. This conclusion contradicted the immediate rush of studies that, without analytical foundation or discipline, promised to deliver on Nixon's promise. However, there was no way the model could be tweaked to produce this answer. Congress and the administration used the PIES model and its successors to analyze changes at the margin, but the policy problem expanded to considering how to operate in a more complicated and more interdependent world. (The Energy Information Administration still produces annual reports using the grandchild of PIES, the National Energy Modeling System, and they can be accessed from *www.eia.doe.gov*.)

3. BEYOND ENERGY INDEPENDENCE

The government's immediate focus in the development and use of the PIES model soon broadened to recognize and exploit many other modeling and analytical efforts to understand the policy choices before the nation and the international energy market. Here there is a vast array of modeling work that deserves attention. I would mention a few strands that are of particular interest to me.

The first has its focus on modeling the behavior of the oil producers in the Organization of Petroleum Exporting Countries (OPEC). The early work by Pindyck developed a dynamic optimization model that characterized OPEC as a monopoly with a competitive fringe and a price elastic demand (Pindyck 1978). The early version of this and studies like it suggested that the potential for further price increases, greater than about \$11 a barrel in 1973 dollars, was limited by the response of demand and that the cartel had already achieved something like the monopoly price. In 1975 there was a major interagency effort within the U.S. government to examine this thesis. The internal analysis showed that the result was highly sensitive to the assumption that oil demand was linear in price. Making the alternative and seemingly innocuous assumption of constant-elasticity demand would produce a completely different result with much higher oil prices for an optimizing cartel. In the absence of good information to bound the second derivatives of the demand curve, the decision was made to classify the conclusion on the grounds that it would not be seemly for the U.S. government to publish this modeling result for the benefit of the oil cartel. Quietly, the next round of government analysis included higher ranges for the assumed scenarios of oil prices (Federal Energy Administration 1976).

Work continued to try to untangle this problem of modeling the behavior of a cartel. This has never been particularly satisfactory, given the complex dynamics and chaotic political situation in the Persian Gulf (Hogan 1996). However, from a methodological perspective, there is special interest in the game theory applications that extended the dynamic monopoly framework to a dynamic Nash-Cournot oligopoly model (Salant 1976). The results capture more of the features of a complicated game where there are strong incentives to free-ride and this creates pressure for the collapse of a restrictive oil production regime.

A key problem in modeling the world oil market is dealing with the uncertainty that prevails and the absence of any simple rational model that explains behavior in traditional economic terms. Revolution and war don't fit well in the optimizing framework. One implication for policy is the need not only to change the level of dependence on uncertain energy markets, but also to develop policies to mitigate the effects of supply interruptions. In the original PI report, the one policy that was accepted without much controversy was the need for a strategic petroleum reserve (SPR) to provide a cushion during an oil supply interruption and the attendant price spike. Once this notion was accepted, the problem was to decide on the optimal size and the policy for using an SPR. The best of the modeling work employed a simplified model of the oil market combined with a simplified Markovian model of uncertainty to produce a not-so-simple and subtle stochastic dynamic optimization problem (Teisberg 1981). The extensive application and augmentation of this model produced two basic insights. First, the optimal size of the SPR was large, larger than any that has been authorized, so we had not made the mistake of making the SPR bigger than the analysis could justify. Second, the benefits of the SPR come from early use, balancing the initial gains from a rapid drawdown to lower the price shock with the longer-term advantage of saving part of the SPR in case the emergency expands. Here the policy record is dismal. In the case of the invasion of Kuwait in 1990, for example, we had one of the largest military mobilizations in the history of the world for operation Desert Storm but we could not declare an emergency and use the SPR to mitigate the attendant price shock. The lesson from the modeling analysis did not translate into policy.

At the time of the focus on PI, there was a parallel effort at the new ERDA, later folded into the Department of Energy, to develop tools for evaluating new energy technologies. The work at Brookhaven National Laboratory used something like the supply-side engineering linear-programming model with a greater focus on technology innovation but no explicit representation of the demand response (Hoffman 1973). In an unrelated effort that blossomed into an extensive series of technology studies, Gulf Oil sponsored development of a network simulation modeling methodology incorporating some of the spirit of the system dynamics model behind *Limits to Growth*. But now there was explicit consideration of prices, demand responses, technology change, and substitution (Cazalet 1975). The resulting SRI-Gulf model was more of a modeling language than a model itself. It could be and was used in many applications that required technology evaluations over long time periods and considering a vast range of substitution possibilities, by organizations like the Electric Power Research Institute and the Gas Research Institute.

The analysis of technology investments over long periods presents a challenge in the evaluation of uncertainty. An early marriage of energy technology modeling using linear programming and decision analysis with an explicit consideration of uncertainty appeared in the extensive evaluations of the breeder reactor (Manne and Richels 1978). If the nation were about to run out of nuclear fuel, large development of the breeder might be warranted. If the nation had adequate supplies of nuclear fuel, no breeder would be needed. But if we didn't know, then a sustained research strategy ("learn, then act") would be indicated. The same decision analytic logic for ex ante optimal policies that are never optimal ex post would influence much of the thinking about research investment. It could have done even more if the framers of the Energy Security Act of 1980 had thought twice about wasting money on synthetic fuels development. An earlier decision analysis modeling effort built on the SRI-Gulf approach has concluded that no subsidy for synthetic fuels investment could be justified on a cost-benefit basis (Synfuels Interagency Task Force 1975). Nevertheless, the United States later launched the Synthetic Fuels Corporation, at great cost and no benefit. This policy failure did not deter further development of the modeling approach that would reappear later in Manne and Richels' (1992) influential work on global warming. See also Nordhaus (1993).

At the time of the Arab oil embargo, the Ford Foundation had been working on a major energy policy study that had a substantial impact on the energy debate. Headed by David Freeman, the study report made the first comprehensive case for the role of energy conservation and increased energy efficiency (Energy Policy Project 1974). Less noticed at the time, but just as important in the long run, was the modeling work that it sponsored in the development of an integrated energy-economic model that described the general equilibrium of the economy as a whole (Hudson and Jorgenson 1974). This was path-breaking econometric modeling that incorporated explicit interactions in a detailed sectoral representation of endogenous input-output coefficients for the economy, and provided intertemporal equilibrium through dynamic investment and the price of capital. The initial work was also married to a detailed engineering model of the energy sector to perform technology evaluations (Hoffman and Jorgenson 1977). Subsequent versions have incorporated endogenous technical change and have been applied to everything from tax policy reform to analysis of the impacts of policies to control climate change (Jorgenson 1998). Here modeling work launched by the energy crisis has gone well beyond to incorporate energy model detail in broader economic policy analysis.

4. CONCLUSION

The work of energy modeling on a large scale built on the early foundations of operations research and economics. The idiosyncratic highlights in the story above cannot give credit everywhere it is due. The purpose rather is to illustrate how the formal analysis and tools of operations research, tools as simple as addition or as complicated as dynamic stochastic optimization, interacted with the policy debate. Energy modeling for policy studies has been a process focused on solving real problems. There has been substantial methodological development, but always as the handmaiden of policy analysis. The work continues today in policy debates that cover policy problems as immediate as electricity restructuring and as long term as global warming. At the turn of the millenium, energy modeling for policy studies was older, but still maturing.

ACKNOWLEDGMENTS

The author is Lucius N. Littauer Professor of Public Policy and Administration, John F. Kennedy School of Government, Harvard University. He thanks L. M. Wein and F. H. Murphy for helpful comments in preparing the paper. The author is solely responsible for any remaining errors.

REFERENCES

Ahn, B. H., W. W. Hogan. 1982. On convergence of the PIES algorithm for computing equilibria. *Oper. Res.* **30** 281–300.

Baumgartner, T., A. Midttun, eds. 1987. *The Politics of Energy Forecasting: A Comparative Study of Energy Forecasting in Western Europe and North America*. Clarendon Press, Oxford, U.K.

Cazalet, E. 1975. SRI-Gulf energy model: Overview of methodology. Working paper. Stanford Research Institute, Menlo Park, CA.

Charnes, A., W. W. Cooper, B. Mellon. 1952. Blending aviation gasolines—A study in programming interdependent activities in an integrated oil company. *Econometrica* **20** 135–159.

Dupree, W., J. A. West. 1972. United States energy through the year 2000. U.S. Department of Interior, Washington, D.C.

Energy Modeling Forum. Stanford University. ⟨www.stanford.edu/group/EMF⟩.

Energy Policy Project of the Ford Foundation. 1974. *A Time to Choose: America's Energy Future*. Ballinger Publishing Company, Cambridge, MA.

Federal Energy Administration. 1974. Project independence report. Washington, D.C.

——. 1976. National energy outlook. Washington, D.C.

Gabriel, S. A., A. S. Kydes, P. Whitman. 2001. The national energy modeling system: A large scale energy-economic equilibrium model. *Oper. Res.* **49** 14–25.

Greenberger, M. 1983. *Caught Unawares: The Energy Decade in Retrospect*. Ballinger, Cambridge, MA.

Harker, P. T., J-S. Pang. 1990. Finite-dimensional variational inequality and nonlinear complementarity problems: A survey of theory, algorithms and applications. *Math. Programming* **48** 162.

Hoffman, K. C. 1973. A unified framework for energy system planning, M. F. Searl, ed. *Energy Modeling*. Resources for the Future, Washington, D.C.

——, D. W. Jorgenson. 1977. Economic and technological models for evaluation of energy policy. *Bell J. Econom.* **8** 444–466.

Hogan, W. W. 1975. Energy policy models for project independence. *Computers and Operations Research*, Vol. 2. Pergamon Press, New York.

——. 1996. Predictions, prescriptions, and policy: Lessons from the energy record. David Lewis Feldman, ed. *The Energy Crisis: Unresolved Issues and Enduring Legacies*. Johns Hopkins University Press, Baltimore, MD.

Hudson, E. A., D. W. Jorgenson. 1974. U.S. energy policy and economic growth. *Bell J. Econom. Management Sci.* **5** 461–514.

Jorgenson, D. W. 1998. *Energy, the Environment, and Economic Growth*. MIT Press, Cambridge, MA.

MacAvoy, P. W., R. S. Pindyck. 1973. Alternative regulatory policies for dealing with the natural gas shortage. *Bell J. Econom. Management Sci.* **4** 454–498.

Manne, A. S. 1958. A linear programming model of the US petroleum refining industry, *Econometrica* **26** 67–106.

——, R. G. Richels. 1978. A decision analysis of the U.S. breeder reactor program. *Energy* **3** 747–767.

——, ——. 1992. *Buying Greenhouse Insurance: The Economic Costs of CO2 Emission Limits*. The MIT Press, Cambridge, MA.

Meadows, D. H., D. L. Meadows, J. Randers, W. W. Behrens. 1972. *Limits to Growth*. Potomac Associates, Washington, D.C.

Murphy, F. H., J. Conti, R. Sanders, S. Shaw. 1988. Modeling and forecasting energy markets with the intermediate future forecasting system. *Oper. Res.* **36** 406–420.

National Petroleum Council. 1972. U. S. energy outlook: A report to the National Petroleum Council's Committee on the U.S. energy outlook. Washington, D.C., December.

Nordhaus, W. 1993. An optimal transition path for controlling greenhouse gases. *Science* **258** 1315–1319.

Pindyck, R. S. 1978. Gains to producers from the cartelization of exhaustible resources. *Rev. Econom. Statist.* **60** 238–251.

Ray, D. L. 1973. The nation's energy future, A report to President Richard M. Nixon. U.S. Atomic Energy Commission, Washington, D.C., December.

Salant, S. W. 1976. Exhaustible resources and industrial structure: A Nash-Cournot approach to the world oil market. *J. Political Economy* **84** 1079–1094.

Samuelson, P. A. 1952. Spatial price equilibrium and linear programming, *Amer. Econom. Rev.* **42** 283–303.

Synfuels Interagency Task Force, to the President's Energy Resources Council. 1975. Recommendations for synthetic fuels commercialization program. Washington, D.C.

Takayama, T., G. G. Judge. 1971. *Spatial and Temporal Price and Allocation Models*. North-Holland, Amsterdam, The Netherlands.

Teisberg, T. J. 1981. A dynamic programming model of the U.S. strategic petroleum reserve. *Bell J. Econom.* **12** 526–546.

U.S. Department of Interior. 1973. The impact of oil interruptions on U.S. energy usage. Washington, D.C.