Rand Journal of Economics
Vol. 16, No. 1, Spring 1985

The effects of technological change, experience, and environmental regulation on the construction cost of coal-burning generating units

Paul L. Joskow*

and

Nancy L. Rose*

This article analyzes the technological, regulatory, and organizational factors that have influenced the costs of building coal-burning steam-electric generating units over the past twenty years. Estimates of economies of scale in construction costs, learning effects associated with utilities and architect-engineers, the costs of environmental regulation, patterns of construction productivity, and cost differences between generating technologies are presented. The results suggest that there are substantial economies of scale associated with the construction of generating units as well as experience effects for both utilities and architect-engineers. The importance of each effect varies across generating technologies. The real costs of generating units have increased dramatically since the late 1960s. These cost increases are only partially attributable to easily measurable responses to environmental restrictions.

1. Introduction

■ In this article we study the technological, regulatory, and organizational factors that have influenced the costs of building coal-burning steam-electric generating units between 1960 and 1980. We are most interested in estimating the effects on construction costs of generating unit size, differences in technology, tightened environmental restrictions, and experience (learning-by-doing) by utilities and architect-engineers. We exploit a data base consisting of 411 generating units built between those twenty years to estimate these effects.

Much has been written on the economics of electricity generation.[1] This article concentrates on one segment of the literature—the capital costs of coal-burning generating

---

* Massachusetts Institute of Technology.

We are grateful for support from the U.S. Department of Energy, the MIT Center for Energy Policy Research, and the MIT Department of Economics. We have benefitted greatly from discussions with Dick Schmalensee, Jerry Hausman, Jim Poterba, and Lew Perl. Many people provided us with comments on earlier drafts. We are especially grateful to Al Klevorick and two anonymous referees for their detailed comments. We want to thank Harry Foster for assistance with data collection and Lew Perl for making his data available to us. The material contained in this article is the sole responsibility of the authors.

[1] The literature is surveyed in Cowing and Smith (1978).

units.[2] These facilities account for over 50% of the electricity produced in the United States. Our work extends the previous literature in a number of ways. First, we control for several "generations" of electric generating technology and focus particularly on differences between the "old" subcritical technology and the "new" supercritical technology. Second, we examine the effects of environmental regulations, which proliferated during the 1970s. Most of the previous literature uses data that predate this period. Third, we recognize explicitly that power plants are not standardized pieces of equipment manufactured in factories, but are brought into operation as a consequence of large-scale construction projects. In this regard, we investigate experience or learning-by-doing effects for firms involved with the design and construction of generating units, and examine the extent to which these effects vary across technologies. Finally, we have assembled and used a more comprehensive and current data base than has been used in previous studies.

The results imply quantitatively and statistically significant scale economies associated with the costs of building coal units. Further, these economies are technology-specific; the more advanced supercritical technology exhibits significantly larger scale economies than does the older subcritical technology. We find evidence of learning-by-doing by both architect-engineers and utilities. These effects also vary across technologies, with supercritical units' exhibiting large and statistically significant experience effects. Controlling for these characteristics, we find that the real costs of building a given coal unit increased by about 100% between the late 1960s and 1980. Part of the increase is associated with scrubbers and cooling towers required to meet new environmental regulations, but most is a "residual" that cannot be explained by easily identifiable responses to these new constraints. The incremental costs probably reflect, in part, unmeasured costs of meeting environmental restrictions, although the cost increases occur prematurely and are too large to be attributable entirely to environmental regulations.

The estimated construction cost relationships reveal a puzzle. At a large scale, the more advanced supercritical technology is both less expensive to build and theoretically more fuel efficient than the best alternative. Yet by the early 1980s, utilities had almost ceased building units using this technology. Our analysis suggests that poor reliability and high maintenance and replacement power costs have made the overall economics of this technology (and probably of very large units regardless of technology) unattractive. The concluding section of the article presents evidence to support this hypothesis.

The article proceeds as follows. The first section describes steam-electric generating technology and its evolution. Next we review the primary determinants of coal unit construction costs and discuss necessary adjustments to utility accounting cost data. We then present the empirical specification of the construction cost model and our estimation methods, and report econometric results for several variants of the basic model. On the basis of these results, we examine how the costs of generating units vary with unit size, across technology, in response to environmental control equipment, as a function of design and construction experience and over time. A discussion of the apparent inconsistency between the observed costs of building generating units and recent utility investment behavior concludes the analysis.

2. Technological and scale characteristics of steam turbine generating technology, 1950–1982

■ The technology for transforming fossil fuel into electricity using steam turbines is fairly old and the basic thermodynamic properties of the steam-turbine (Rankine) cycle

[2] We view this work as a natural extension to coal generation of Zimmerman's (1982) analysis of nuclear plant costs. The article also extends that portion of Wills (1978) which examines fossil plant construction costs for the period 1958–1970 and Perl's (1982) unpublished work. Finally, the article helps to fill a gap in Bushe's (1981) study. In this otherwise excellent study of generating technology, essentially no analysis of the actual capital costs of different types of units is provided.

are well understood. Fossil fuel[3] (usually coal) is burned in a furnace to generate pressurized high temperature steam (the furnace and the steam-generating and heating equipment are generally referred to collectively as "the boiler"). The pressurized steam is expanded through a turbine which turns a generator to produce electricity. The steam exhausted from the turbine is then cooled in a condenser and returned to the boiler to begin the cycle once again. Central station generating units using this cycle were first introduced in the United States around the turn of this century, and have since been the primary basis for electricity generation in the United States.[4]

The overall thermal efficiency of a steam generation cycle increases with the temperature and pressure of the steam, the thermal efficiency of the boiler, the efficiency of the turbine, and the size of the boiler and turbine. Over time, a large number of technological advances have increased design thermal efficiencies.[5] By the late 1950s, however, a steam temperature threshold of sorts in the 1000°F to 1010°F range was reached; almost all new units installed since then have this steam temperature.[6] Many other design changes that historically led to large increases in thermal efficiency had become standard by the early 1960s.[7] As a result, since 1960 the primary technological frontier on the thermal efficiency front has been in the pressure and, to a lesser extent, the size dimensions.[8]

Table 1 reports the steam pressure characteristics of all coal units installed between 1950 and 1982, measured as a percentage of new coal capacity placed in operation during each time period.[9] Generating units are divided into five groups on the basis of rated turbine throttle pressures. Units constructed after 1959 fall naturally into at least three pressure classes grouped around 1800, 2400, and 3500 psi, with a possible fourth class grouped around 2000 psi.[10] These four pressure classes aggregate into two major technological classes: *subcritical* units with pressures less than 3206 psi and *supercritical*

---

[3] Coal accounted for 53% of total electricity generation and 72% of fossil-fuel steam generation in 1982. Nuclear plants are built around the same basic steam cycle, but rely on the heat from nuclear fission to produce steam. Nuclear plants operate at lower steam temperatures and pressures than state-of-the-art fossil plants.

[4] See, generally, Ling (1964), Bushe (1981), Cootner and Lof (1965, chaps. 3, 5, and Appendix A), *Mark's Standard Handbook for Mechanical Engineers* (8th ed., pp. 9-54 to 9-56), and *Electrical World* (June 1, 1974, pp. 78–81).

[5] For a more detailed discussion of steam-electric generating technology, see Ling (1964), Cowing (1974), Cootner and Lof, (1965, chaps. 3, 5, and Appendix A), *Mark's Standard Handbook for Mechanical Engineers* (8th ed., pp. 9-54 to 9-56), and Bushe (1981).

[6] There is no theoretical reason why higher steam temperatures cannot be achieved, but technical and economic constraints appear to have made 1000° to 1050°F the maximum practical temperature for a Rankine steam cycle during the past two decades. The first unit in the 2400 psi/1000°F range became operational in 1953. The first supercritical unit became operational in 1957 with a steam pressure of 4500 psi and steam temperature of 1150°F. A subsequent unit installed in 1960 had a steam pressure of 5000 psi and steam temperature of 1200°F. Despite these early experiments with very high pressure and temperature steam in supercritical units, almost all supercritical units installed since 1960 have had steam conditions close to 3500 psi/1005°F. See *Electrical World* (June 1, 1974, pp. 78–79) and Cootner and Lof (1965, Appendix A). See also *EPRI Journal* (January/February 1984, p. 46).

[7] These include the introduction of one or sometimes two reheat cycles and preheat cycles using multiple bleed points from the turbine to increase average cycle efficiency. See Bushe (1981) for a detailed discussion.

[8] Improvements in design efficiency due to increases in unit size appear to be relatively small for central station units of the sizes that make up the bulk of our sample. An increase from 350 to 700 Mwe electric increases design efficiency by .5% (Sargent and Lundy, 1980, p. 3). See also *Mark's Standard Handbook for Mechanical Engineers* (8th ed., pp. 9-54 to 9-56) and Bushe (1981, chap. 3) regarding the basic technology.

[9] The sample used for econometric analysis includes 411 units built between 1960 and 1980. We do not have data on all of the necessary variables for a longer period of time.

[10] We chose this grouping after reviewing the engineering literature (e.g., *Mark's Standard Handbook for Mechanical Engineers,* 8th ed., pp. 9-54 to 9-56) and examining the distribution of units by rated turbine throttle pressure. The units cluster tightly at four pressure nodes: 1450, 1800, 2400, and 3500 psi. In addition, there are a few units at 2000 psi. We were unable to find much discussion of the 2000 psi group in the literature, although Bushe (1981) lists it as a separate category. Rather than arbitrarily allocating these units to the 1800 or 2400 psi classes, we treat them as a distinct group.

TABLE 1 **Capacity Additions by Technological Group and Year: 1950–1982 (% of New Capacity)**

Period	Turbine Throttle Pressure Groups (psi)				
	Subcritical				Supercritical
	1600 or less	1800	2000	2400	3500
1950–1954	39	45	13	2	0
1955–1959	10	32	36	20	1
1960–1964	2	21	20	45	12
1965–1969	2	8	1	46	42
1970–1974	0+	5	0+	32	62
1975–1980	0	6	1	62	31
1981–1982	0	2	4	88	6

Source: See text.

units with pressures greater than 3206 psi. We discuss the technological and economic differences between subcritical and supercritical units in the next section.

There are three things to note in Table 1. First, until the mid-1970s we see a pervasive movement of the industry from lower pressure units to higher pressure units. This is a continuation of a trend that began with the first steam turbine units introduced at the turn of the century (*Electrical World,* June 1, 1974, p. 79; Cowing, 1974, p. 147). Second, the new technologies with higher pressures and associated higher design efficiencies only gradually replace older technologies, so that different technologies coexist at any point in time. Third, the "inevitable" movement to the higher pressure supercritical units began in the early 1960s, continued into the 1970s, and then reversed itself. Today, supercritical technology has almost been abandoned.

The movement to higher pressure units was accompanied by a fairly rapid increase in the average size (measured by generating capacity) of new units until about 1975. Table 2 presents data on the mean, minimum, and maximum sizes of coal units placed in operation between 1950 and 1982. The trend towards increasing size leveled off by the late 1970s; the average size of units installed has declined slightly since 1975. Very large units (above 750 Mwe) are no longer being built by utilities.

3. Primary factors affecting the construction costs of coal generating units

■ A steam generation and electricity production system is created as the result of a large and complicated construction project that incorporates many unit-specific design characteristics and requires a considerable amount of time to complete. It involves a

TABLE 2 **Size Distribution of New Coal Capacity by Year: 1950–1982 (Mwe)**

Period	Mean	Minimum	Maximum	Number of New Units Installed
1950–1954	124	100	175	99
1955–1959	168	100	335	175
1960–1964	242	100	704	104
1965–1969	407	103	950	100
1970–1974	591	115	1300	109
1975–1980	545	114	1300	127
1981–1982	517	110	891	41

Source: See text.

detailed design and engineering process, the procurement of a large number of components from many different suppliers, the assembly of these components into a steam generation and electricity production system on site, and the construction of structures to house the assembled components. The construction requires a large labor force composed of workers from many different crafts (Willenbrock and Thomas, 1980, p. 799). Actual construction typically takes three to five years,[11] and the magnitude and complexity of the project require extensive design and engineering work and specialized managerial skill for it to be accomplished at minimum cost.

Cowing (1974, pp. 140–141) identifies four primary factors that determine the nominal costs of a generating unit (although he does not estimate a construction cost function for generating units directly).[12] These are: the unit's size, its thermal efficiency, its "vintage," and input (equipment) prices. We incorporate these factors, as well as several others, in the construction cost function that we estimate below. The causal factors that are of most interest to us are unit size, choice of generating technology as characterized by the design steam pressure (an *ex ante* measure of thermal efficiency), construction productivity, the effects of compliance with environmental regulations, and the contribution of experience or learning-by-doing at the construction stage. We discuss each in turn.

□ **Size: economies of scale in construction.** The engineering literature generally assumes that capital costs increase less than proportionately with the capacity of the unit. For example, a recent engineering cost analysis comparing 350 Mwe with 700 Mwe coal units suggests that average construction cost per unit of capacity declines by about 20% when the unit size is doubled (French and Haddad, 1980; Electric Power Research Institute, 1982, Appendix B, p. 55; Dow Chemical Company, 1975; Oak Ridge National Laboratory, 1972). Early econometric studies using data on relatively small units built primarily in the 1950s found similar qualitative results (Ling, 1964, pp. 74–77). But more recent studies have questioned these findings. Wills (1978, p. 505) found minimal economies of scale in construction costs beyond 100 Mwe of capacity (the minimum size unit in our sample) for units built between 1958 and 1970. Stewart (1979, p. 559) found diseconomies of scale over much of the size range in his sample.[13] Zimmerman (1982) found that cost estimates for nuclear plants were based on a scale parameter twice the size of the realized parameter.

Our data set includes more recent data and much larger units than previous studies, as well as a range of unit sizes that varies by an order of magnitude from smallest to largest. This should allow us to test whether economies of scale continue to be observed for larger, more recent units, as suggested by the engineering literature and most earlier econometric studies.[14]

□ **Differences in steam pressure and design thermal efficiency.** As we discussed above, technological change in this industry since the early 1960s has focused on advancing the

[11] This reflects only the actual time to completion once construction is started. It does not include time required for design and engineering work, regulatory approvals for construction permits, etc. The "planning/construction" cycle may be as much as eight years today. See "Delays and Cancellations of Coal-Fired Generating Capacity," Energy Information Administration, U.S. Department of Energy (July 1983, p. 19) and Electric Power Research Institute (1982, Appendix B, p. 55).

[12] Cowing does not estimate the construction cost function directly. He uses a particular analytical specification of such a function to derive reduced-form relationships for the optimal values of capacity and efficiency, which depend on both fuel and capital costs.

[13] This result is probably a consequence of the inappropriate pooling of combustion turbine units and steam cycle units. He also has observations on only 19 steam units.

[14] The extent to which there are scale economies in generation has important implications for the optimal structure of electricity-producing firms. See Joskow and Schmalensee (1983, chap. 7).

pressure and size frontiers. For subcritical units, increasing steam pressures required the development of stronger materials to withstand the higher pressures, as well as thicker casings for components (Bushe, 1981, p. 47). Aside from additional materials' costs to accommodate higher pressures, however, there were no major technological differences among the subcritical pressure groups during our sample period.[15] Thus, we expect that cost differences among subcritical technologies were probably fairly small by the 1960s, with higher pressure units having somewhat higher costs than lower pressure units, other things equal.

The development of supercritical boilers represented a more fundamental departure from previous technology. Water heated to a temperature above 706°F at a pressure above 3206 psi directly vaporizes to dry superheated steam. This eliminates equipment required to extract saturated steam, recycling equipment, and some equipment to heat saturated steam, but requires additional expenditures on materials to accommodate the large increase in steam pressures. Engineering cost calculations suggest that supercritical units are characterized by substantial economies of scale in construction. Recent engineering calculations suggest that supercritical units may be less expensive to build than state-of-the-art subcritical units at scales above 500 Mwe (Electric Power Research Institute, 1982, Appendix B, p. 55).

□ **Real cost changes over time: construction productivity and compliance with environmental regulations.** Two primary factors are likely to have increased the real costs of building coal generating units (after accounting for input price changes) over time. First, because a coal unit is a major construction project, its costs are sensitive to changes in construction productivity. Analysis of the construction costs of nuclear units suggests that the number of labor hours per unit of capacity has increased dramatically over time.[16] A large fraction of this increase often is attributed to changes in nuclear plant safety and environmental regulations, which entail more equipment, construction delays, and reconstruction of portions of some plants. But utility managers to whom we have spoken suggest that even given these changes, construction productivity has generally deteriorated over the past decade. Zimmerman (1982) finds a large and significant trend in construction costs of nuclear units, which implies rising real costs over time. We wish to determine whether similar secular changes in costs are found in the construction of coal units.

Second, we expect to observe that more stringent environmental requirements increased the cost of building coal-burning units. Restrictions on power plant discharges of effluents into the environment tightened dramatically during the 1960–1980 period we study. These are a consequence of environmental legislation—and associated administrative regulations promulgated as a result of this legislation—enacted primarily during the 1970s. The legislation includes amendments to the Clean Air Act, the Federal Water Pollution Control Act, the Solid Waste Disposal Act, the Noise Control Act, and related state laws.[17]

---

[15] Single reheat cycles and 5 to 7 bleedpoints for boiler feedwater preheating had become fairly standard in all units, regardless of pressure, by about 1960.

[16] This was emphasized in the Department of Energy's comments on an earlier draft. See also *Nuclear Plant Cancellations: Causes, Costs, and Consequences,* U.S. Department of Energy, DOE/EIA-0392, (April 1983).

[17] The Clean Air Act Amendments of 1970 and 1977 (PL91-604 and PL95-95) and the Water Pollution Control Act Amendments of 1972 (PL92-500) had the most direct and probably the largest impact on electric generating plants. See 40 CFR 60 and 40 CFR 423 for the associated regulations. In addition, solid waste storage and disposal (ash, slag, sludge from scrubbers, etc.) is subject to regulation under the Solid Waste Disposal Act of 1976 (PL96-580), and power plant noise is regulated under the Noise Control Act of 1972 (PL92-574). See note 18. Furthermore, since power plants involve a substantial amount of on-site construction, construction practices are subject to regulation under the Occupational Safety and Health Act of 1970. There were 200 pages of safety and health regulations for construction by 1979 (29 CFR 1926).

Unfortunately, although we can determine which units were built with scrubbers and cooling towers, we cannot observe all of the design and construction responses to the plethora of new environmental restrictions imposed during this period.[18] We also cannot account separately for other factors that may have affected construction productivity over time. After controlling for the larger, easily identifiable construction responses to environmental constraints, we expect to observe a "residual" secular increase in the real costs of building coal plants at least during the 1970s. This will reflect "unmeasured" environmental control costs, as well as the effects of changes in construction productivity.

We account for changes in real costs over time by estimating yearly "time effects." A dummy variable is introduced for each year, 1960 through 1979, and it is set equal to one if the unit entered commercial operation during that year and zero otherwise. We emphasize that the time effects will pick up "unmeasured" costs of environmental compliance *as well as* other factors that have affected construction productivity. The independent contributions of each cannot be identified.

□ **Learning and experience effects.** At least two studies of nuclear power plant construction costs have identified learning or experience effects (Zimmerman, 1982; Mooz, 1978). They have found that as the cumulative number of units built increases, other things equal, the average cost per unit of capacity declines. These learning effects typically are identified with the architect-engineer or constructor/construction manager (often the same firm for a nuclear plant). We investigate the presence of these learning effects in coal unit construction.

The process of building a coal unit starts with the design and engineering of the unit. For most utilities this task is assigned to an architect-engineer, who prepares initial engineering plans and cost estimates and continues to work on design and engineering problems through the construction phase of the plant. Some utilities handle design and engineering internally, although the number of utilities that do so has declined over time as construction projects have become larger and more complex (Willenbrock and Thomas, 1980, chaps. 6–8).

What then are the sources of any learning effects that might be observed? The traditional sources of learning in airframe and shipbuilding construction associated with "direct learning by labor" (Searle and Goody, 1943; Alchian, 1963; Hirsch, 1952) seem especially unlikely in the actual construction of plants, given the nature of the construction process and the use of labor. Equipment manufacturers may experience traditional learning economies in factory production of components. But to the extent that these economies are reflected in declining procurement costs, if the construction cost data are adjusted for input price changes (as we adjust them), these learning effects should not be observed in the real construction cost data.

At the actual construction stage, learning effects seem most likely to be associated with the repetitive design of technologically similar power plants and with repetitive management of construction. Some of this learning may accrue to individual architect-engineers, construction managers, and utilities, but "general" or "industrywide" appropriation of design innovations is also possible. Although utilities frequently retain other firms as their agents to do design and engineering work and to manage construction, they are generally intimately involved with the whole process. Utilities consequently may acquire experience that can mitigate potential principal-agent problems, and they may

[18] There are numerous "small" environmental, health, and safety related expenditures that were required by utilities over this period. See for example, Electric Power Research Institute, "Noise Control at Fossil Fuel Power Plants: An Industrywide Assessment of Costs and Benefits," EPRI-CS-3262 (December 1983). Since 1960, utilities spent more than $200 million for the control of noise from fossil plants, at an increasing rate of expenditure over time.

become especially knowledgeable about local labor market conditions and experienced in dealing with local craft unions. Therefore, both architect-engineer and utility experience in the power plant design and construction process could help to reduce costs.

A number of difficulties are generally associated with the measurement of experience effects (Gold, 1981, pp. 17–18). The first is the problem of distinguishing cost savings associated with experience from those associated with technological change more generally and with exploitation of scale economies. Second, there are questions about the appropriate way to measure experience. Do we count experience with broad classes of products or processes (e.g., coal plants generally) or do we measure experience within narrower classes (e.g., subcritical vs. supercritical units)? Third, any simple measure of experience represents at best a characterization of a very complicated process that is not yet well understood theoretically. What is the appropriate functional form for experience variables? Do experience effects persist over long periods of time or do they deteriorate over time?

We cannot provide definitive answers to all of these questions. Learning-by-doing is one aspect of technological change which seems in general to be poorly understood by economists. We consider a number of approaches to measuring experience effects and believe that our analysis substantially improves upon previous work.

Following Lieberman (1984), we attempt to distinguish among scale economies, technological change, and experience effects by introducing separate variables for unit size, technology, time, and experience into the statistical analysis. Further, we build considerable flexibility into the analysis. Cost levels and scale parameters are allowed to vary across the four technology classes; independent time effects are estimated to capture (among other things) the traditional "residual" measure of technological change. Experience is calculated separately for architect-engineers and utilities, in terms of both total experience and technology-specific experience. We consider two functional forms for the experience measures. We allow experience coefficients to vary across technologies, and test for potential interactions of architect-engineer and utility experience effects. Finally, we test whether experience effects are constant over time.

4. Data, accounting issues, and adjustments for input price changes and regional effects

■ Our data base consists of 411 coal units, or about 95% of all units greater than 100 Mwe designed to burn coal, that began operating between 1960 and 1980.[19] The coal units missing from our sample were excluded because of missing data, usually missing construction cost data. Data sources and transformations are described in the Appendix. We discuss below several difficulties with reported data.

Although our primary sample period for construction costs covers the period 1960 to 1980, units in each pressure class were also built before 1960. Information on units built before 1960 is therefore necessary to construct architect-engineer and utility

[19] The units fall into the following pressure and size classes:

Pressure (psi)	# of Units	Size Range (Mwe)		
		Min.	Max.	Mean
1800	82	100	617	191
2000	22	136	598	291
2400	197	136	893	446
3500	110	325	1300	723

experience measures. To obtain this information, we performed a census of all generating units placed in operation since 1949 and assembled information on their turbine pressures, utility operators, and architect-engineers, by using a variety of engineering publications and interviewing plant operators when necessary. We chose 1949 as the cutoff because Bushe's data (1981) indicate that relatively few units in the classes represented in our primary sample were constructed before 1950, and because the data collection effort becomes increasingly difficult as the cutoff date is pushed back. We believe that we have an essentially complete experience history of units with pressures of 2000 psi and above and a reasonably complete history for the 1800 psi group.

Although our interest is in the determination of real construction costs, utilities report only the nominal book costs of construction associated with each generating unit. Three potential problems emerge. The first is what we call the "first-unit effect." The majority of the units in our sample are part of multiunit sites where two or more units are planned to be built in sequence over a period of a few years. There are substantial common costs associated with design work, site procurement and preparation, coal handling facilities, water intake and discharge facilities, control rooms, waste storage facilities, etc. Because of utility-commission rate-setting procedures, utilities have strong incentives to assign as large a fraction of these common costs as possible to the first unit built on a site. We therefore expect the accounting data to indicate that first units are more costly than follow-on units, a reflection of this allocation of common costs.[20] Failing to control for this could bias upward our estimates of learning effects, since on average as more units of a particular generating technology are built, they are more likely to be follow-on units than first units.[21]

The second issue involves the appropriate measurement of construction costs. The observed cost data reflect nominal expenditures on construction plus interest charges accumulated during construction.[22] Over time, input prices and interest rates have changed, and it appears that the typical construction time has increased somewhat (Komanoff, 1981; U.S. Department of Energy, 1983). Ideally, we want a measure of real construction costs that deflates nominal construction costs to account for changes in input prices and interest rates. But we want to retain cost differences due to changing construction times since longer construction periods for the same piece of capital equipment represent a real increase in the effective cost of the unit.

Construction times and cash flow profiles are not published for individual coal-generating units. We therefore use information on the typical cash flow pattern for coal units built in the early 1970s, along with a regional input price index and a time series of interest rates, to construct a "standardized" real construction cost variable net of accumulated interest charges.[23] Other things equal, units that took longer to build will have a higher standardized cost, reflecting residual interest charges in the cost figures and *vice versa*. Thus, if construction times have increased, so too should real costs, after accounting for other factors affecting construction costs. If it takes longer to build larger

[20] A recent engineering cost study suggests that first-unit costs are about 20% higher than follow-on unit costs (French and Haddad, 1980, Table 1).

[21] Zimmerman (1982) finds a large statistically significant first-unit effect in his analysis of nuclear plant costs. Wills (1978) does not account for this effect directly, but does distinguish between single-unit and multiunit plants.

[22] The construction cost figures that we obtain from utility accounting data normally include accumulated interest charges (interest during construction) incurred during the construction of the plant. But interest charges are not reported separately from ordinary construction expenditures. Other things equal, interest-during-construction charges increase with nominal interest rates and with construction times.

[23] The procedure is discussed in more detail in the Appendix. The cash flow pattern is from "Powerplant Capital Costs—Current Trends and Sensitivity to Economic Parameters," U.S. Atomic Energy Commission, WASH-1345 (October 1974).

units than smaller units, the scale economy estimates will reflect the associated cost differences. Although this approach is not ideal, it is the best that we can do absent the actual cash flow pattern for each unit.

We use the seven-region Handy-Whitman index for steam-generating construction costs as our input price deflator. We believe that this is the best index available to control for input price changes.[24] To check the sensitivity of the results to this choice of index, all equations below were also estimated by using the raw nominal cost data. Except for the levels of the time effects (but not their general pattern), the results are not affected in any important way.[25]

The third problem arises from regional cost variations. The Handy-Whitman index controls for cost *changes* on a regional basis, but it does not control for differences in the level of costs across regions.[26] Although we are not interested in regional cost differences *per se,* their omission potentially biases some of the coefficients of interest. Regional variations arise from several sources.

First, construction labor is hired largely in local and regional markets, across which wages vary considerably. The extent of unionization and union work rules also varies widely across the country and may lead to systematic regional differences in construction costs. Since three of the four largest utilities that do their own engineering and construction work happen to be located in relatively low wage areas, it is important to account for wage variations. Failure to do so may confound utility-specific cost effects and learning effects with what are simply differences in input prices.

Second, climate conditions affect the cost of materials and construction. In colder climates, enclosed structures are built to house boilers and turbine-generator equipment indoors. In warmer climates, either the boiler or the turbine-generator equipment or both may be placed outdoors, with a saving in construction costs. Variations in overall costs due to differences in structures are likely to be small, however: structures account for only about 15% of total construction costs. Weather extremes (cold, heat, and rain) may also adversely affect construction productivity; their impact may vary across regions. Finally, the type of coal that a power plant is designed to use may affect construction costs. Schmalensee and Joskow (1985) found this to be the case in a study of a much smaller number of units built during the 1960s. Since coal quality varies from region to region (Joskow, 1985), this too may lead to systematic regional differences in construction costs.

We are able to account for differences in structures (indoor/outdoor) and regional wage differences for all units in our sample. We therefore rely on these two variables in the econometric work reported below. But we also report estimates allowing for regional fixed effects to see whether other omitted regional characteristics are likely to impart a serious bias to our results. This potential source of bias does not appear to be a problem.

---

[24] This is a proprietary index that has been made available since 1912. The precise composition of the index is not made public, but it appears to be closer to a fixed-weight price index rather than a variable-weight hedonic price index reflecting changes in design characteristics, productivity, scale, environmental requirements, etc. Indeed, in discussing the index the supplier states explicitly that: "State-of-the-art changes often affect costs independently of inflation. New regulatory and environmental requirements, changes in work rules and improved design standards, for instance, increase construction costs even though the price of wages, materials, and equipment may be static. Trended [using the H-W index] construction costs will not reflect such changes." ("The Handy-Whitman Index of Public Utility Construction Costs," Bulletin No. 118, Baltimore: Whitman, Requardt and Associates, July 1, 1983, p. ix). The fact that the index does not reflect these changes is ideal for our purposes. It should be noted that, although the index captures regional *changes* in input prices over time, it tells us nothing about the relative *levels* of input prices by region. It is therefore necessary to control separately for differences in wages and construction practices across regions.

[25] When we use nominal cost data, the residual time effects include the effects of input price changes. The coefficients of the time dummy variables then yield a hedonic price index for coal-generating units.

[26] The index sets 1949 costs equal to 100 for each region. Escalation rates are computed *within* each region, but there is no way to control for variations *across* regions.

5. Specification of the construction cost model

■ The previous discussion suggests a construction cost function that relates the real cost of a unit to the unit's size, the presence of scrubbers and cooling towers, experience variables, and time effects, along with appropriate adjustments for changes in input prices, accounting peculiarities, and regional effects. Following Zimmerman (1982), we specify the relationship between unit cost (cost/Kw) and size to be Cobb-Douglas, a specification that is also used in many engineering cost models (French and Haddad, 1980; Electric Power Research Institute, 1982; Dow Chemical Company, 1975; Oak Ridge National Laboratory, 1972). This forces the cost function to have a constant elasticity of unit cost with respect to size. But by allowing the cost function to have different coefficients for different technologies, we build additional flexibility into the cost/size specification. As discussed below, we have estimated alternative specifications that allow the elasticity of unit cost with respect to size to vary with size and to change sign so that regions of economies and diseconomies of scale can be observed. The more flexible specification does not improve or alter the results in any important way.

Several variants of the following basic construction cost relationship are estimated. The assumed error structure and the estimating technique are discussed in the next section. The data are described in further detail in the Appendix. The basic construction cost relationship is

$$LAC = \sum_{t=1960}^{1979} a_t T_t + b_1 LSIZE + b_2 RWAGE + b_3 FIRST + b_4 SCRUBBER + b_5 COOLTWR + b_6 UNCONV + b_7 EXPERAE + b_8 EXPERU + b_9 EXPERI + S + u, \quad (1)$$

where

LAC = the natural logarithm of the standardized real cost per Kw of a generating unit expressed in $1980 net of capitalized interest costs.

T_t = one if the associated unit commenced commercial operation in year t and zero otherwise. The dummy variable is omitted for 1980, so that the estimated coefficients are relative to 1980 cost conditions.[27]

LSIZE = the natural logarithm of unit size expressed in thousands of Kw's (megawatts). We estimate the model by constraining the size coefficient to be constant across pressure groups (homogeneous technology) and also by allowing the size coefficient to take on a different value for each pressure group (heterogeneous technology). This allows us to test for differences between technological groups and to see whether a failure to specify technological differences affects the estimated values of certain coefficients.

RWAGE = The regional average union wage for construction workers in 1976 for the Bureau of Labor statistics region in which the unit is located. We have experimented with regional wage data for other years, and they perform just as well as a proxy for regional wage differences.

FIRST = one if the unit is the first unit on the plant site, zero otherwise.

---

[27] We have also estimated these relationships with a linear time trend and with a quadratic time variable. As we shall see, it is quite clear that there are important residual time effects in the data. These effects are not linear as in the case of Zimmerman's (1982) analysis of nuclear plant costs, however. Rather, they are roughly quadratic with real costs per unit of capacity, other things equal, declining during the 1960s and increasing during the 1970s. The use of a quadratic in time rather than year-by-year time dummies does not affect the results in any important way.

SCRUBBER = one if the unit was built with a scrubber, zero otherwise. Retrofit scrubbers are coded as zero since their costs would not be reflected in the initial construction cost data.

COOLTWR = one if the unit was built with any type of cooling tower, zero otherwise.

UNCONV = one if the unit is not completely of indoor design, zero otherwise.

EXPERAE = a measure of the relevant experience since 1950 of the architect-engineer that designed the observation unit. For a unit that enters commercial operation in year *t*, the experience is measured as the cumulative number of "like" units that entered service before year *t*. Thus, for a given architect-engineer, all "like" units placed in operation in a particular year are assigned the same experience value. We define "like units" in a variety of ways, as discussed below.

EXPERU = utility experience since 1950, reflecting either total utility experience with coal units or utility experience within each technological group. Utility experience is also measured in a variety of ways, described below.

EXPERI = a measure of cumulative "industry experience." This follows Zimmerman's work (1982), which includes total industry experience with nuclear plants to measure the "nonappropriable" learning effects that accrue to the industry as a whole in the construction of power plants.[28]

We estimate all equations by using two different functional forms for the experience variables: the conventional log of total experience (actually one plus experience), and Zimmerman's [1/(1 + *experience*)] specification. For the former, experience effects imply a negative coefficient; for the latter, a positive coefficient.

S = an intercept term. As for the size coefficients, we estimate the model both by constraining this to be the same for all pressure groups and by allowing it to vary across pressure groups.

u = an error term, the structure of which we discuss in more detail in the section below.

6. Error structure and estimation technique

■ Our data set consists of 411 generating units designed by many different architect-engineers. We believe that there are likely to be architect-engineer-specific design characteristics common to units designed by a particular firm.[29] We therefore treat the data set as a panel, with individual generating units as observations over time on a cross section of architect-engineers. The data set can be thought of as a cross section times series, except that generating units replace the time index. This substitution is necessary both because the panel is unbalanced, in that there are different numbers of units for different architect-engineers, and because we have multiple units per calendar year for some architect-engineers.

[28] We cannot identify both the year-by-year effects and industry experience effects in the homogeneous technology specification in which we use all-coal experience (rather than technology-specific experience) as the experience measure. Given the way experience is measured (year by year), the industry experience variable in this case is just a linear combination of the time dummy variables. We cannot identify separate industry experience effects for each technology as well as the time dummy variable coefficients for exactly the same reason.

[29] For example, American Electric Power is typically thought to place greater emphasis on plant reliability, and thus perhaps causes higher initial construction costs than those associated with less stringent design standards. Bechtel designs plants using Combustion-Engineering boilers more frequently (but not exclusively) than does the average architect-engineer.

This structure suggests writing our model as:

$$LAC_{ij} = \sum_{t=1960}^{1979} a_t T_{ijt} + X_{ij}B + u_{ij}, \qquad i = 1, \ldots, N, \qquad j = 1, \ldots, J_i, \tag{2}$$

where i indexes architect-engineers, j indexes units, a_t is the time effect for units built in year t, T_t is the dummy variable for year t, X consists of the remaining variables in equation (1), and B is the vector of parameters associated with these variables. The error term has the structure

$$u_{ij} = v_i + h_{ij}, \tag{3}$$

where v_i is an error term specific to architect-engineer i and common to all units designed by this architect-engineer, and h_{ij} is a unit-specific error term. If we make conventional panel data assumptions about the moments of the error components,[30] including the assumption,

$$E(v_i|X) = 0, \tag{4}$$

we can obtain consistent estimates of B in (2) by using ordinary least squares (OLS), generalized least squares (GLS), or fixed effects estimation techniques. Of these, GLS will be efficient.

Equation (4) will not hold if the exogenous variables are correlated with the architect-engineer-specific effects. There is reason to believe that such correlations exist in our data. The v_i represent characteristics of particular architect-engineers, such as design philosophy, choice of vendors for major components, and the quality of the units designed. Because these unmeasured attributes of different architect-engineers may be correlated with experience, failing to account for them could lead to biased estimates of the parameter vector B. For example, if architect-engineers that have designed a relatively large number of units happen to concentrate component orders with a particular boiler or turbine manufacturer, and if these manufacturers in turn charge lower prices than other vendors for equipment, differences in equipment prices may show up as architect-engineer experience effects in the absence of manufacturer-specific input prices, if we do not control for architect-engineer-specific effects. Alternatively, if "high-quality" architect-engineers design more units, and higher quality entails higher costs, the absence of information on quality attributes may prevent us from detecting experience effects when they are present, unless we control for the architect-engineer-specific effect.

These possible correlations affect our choice of estimating technique. If (4) is not valid, then both OLS and GLS estimation of (2) will yield inconsistent estimates of B. Fixed effects—that is, conditioning on the v_i or explicitly estimating a separate intercept for each architect-engineer—will be consistent.[31] Fixed effects estimation is appropriate also when the estimated values of the architect-engineer effects are of interest in their own right. If we can maintain (4) and are not interested in particular architect-engineer effects, then GLS will be preferred.

We test the validity of (4) by a specification test that follows Hausman (1978) and Hausman and Taylor (1981). We can reject the null hypothesis that $E(v_i|X) = 0$ at the .10 level for each of the variants of the model presented below.[32] This cautious rejection

---

[30] See Nerlove (1971) for an exposition of these assumptions.

[31] Hausman and Taylor (1981) describe an instrumental variables estimator that is both consistent with and efficient for this general class of problems. Since we have some interest in actually obtaining estimates of architect-engineer effects directly, for use in future work, we use fixed effects rather than instrumental variables with some sacrifice in efficiency.

[32] We test the validity of (4) by performing a specification test for each of the three variants of equation (1): homogeneous technology, heterogeneous technology with equal learning effects, and heterogeneous technology

and our interest in analyzing the performance of different architect-engineers in future work lead us to use fixed effects estimation for all equations reported below. This yields an estimated intercept for each architect-engineer. We report the mean value of these intercepts for each equation below.

7. Empirical results

■ We report and discuss the econometric results in this section. In Section 8, we examine the quantitative effects of size, thermodynamic efficiency (steam pressure), environmental regulation, learning-by-doing, and residual productivity changes on construction costs. The estimated parameter values are presented in Table 3.

□ **Homogeneous technology.** The simplest specification of the model ignores the likely differences in technologies and treats all coal units as drawn from a homogeneous technology. This involves estimating equation (1) while constraining the coefficients to be the same for all technologies. This specification requires implicitly that we measure experience for architect-engineers and utilities in terms of *total* coal unit experience rather than technology-specific experience. The coefficient estimates and standard errors for this specification using the log of cumulative experience as the experience measure are reported in column 1 of Table 3.[33] We report the coefficient of the time dummy variable for 1970 for this and all subsequent equations, and postpone discussion of the pattern of the residual time effects to the next section.

We obtain a very precise estimate of the scale coefficient, equal to −.183. This is quite close to the value Zimmerman (1982) found for nuclear plants. The very large and precise estimate of the first-unit effect implies that first units' accounting costs are 25% higher than those of follow-on units. When this effect is allowed to vary with the final number of units built on the site, the estimated coefficients increase with the number of units, but are not significantly different from one another. The regional wage effect is fairly large with a small standard error. At the mean wage rate, the elasticity of construction costs with respect to the wage rate is about .35. As expected, the use of outdoor construction reduces costs, although the effect is fairly small and imprecisely estimated. Scrubbers add about 15% to the cost of the plant, and the effect is statistically

---

with variable learning effects. The procedure follows Hausman (1978) and Hausman and Taylor (1981). The test is based on the estimate of

$$LAC_{ij} = \sum_{t=1960}^{1979} a_t T_{ijt} + (X_{ij} - X_i)B + (X_i)G + u_{ij}, \tag{i}$$

where $[T_t]$ is the set of dummy variables from equation (1), X_{ij} is the vector of values for unit ij for all other variables as described in equation (1), X_i is the vector of architect-engineer means for these variables,

$$X_i = (1/J_i) \sum_j X_{ij}, \tag{ii}$$

$[a_t]$, B, and G are the associated parameter vectors, and $u_{ij} = v_i + h_{ij}$. Under the null hypothesis of no misspecification, the two estimated vectors $\hat{B}$ and $\hat{G}$ should be identical. If, however, $E(v_i|X) \neq 0$, we should reject $\hat{B} = \hat{G}$.

Note that this test maintains the conventional panel data assumption that the time effects are independent of the architect-engineer-specific fixed effects. Although we believe experience measures are the most likely source of possible correlation of the independent variables with the architect-engineer-specific effects, we test the entire X vector. Comparison of the fixed effects and OLS point estimates suggests that the experience coefficients are the primary cause of our (cautious) rejection of (4). Note also that this is one of several ways to construct a test of assumption (4).

[33] Using (1/1 + *experience*) has no effect on the results reported for this specification, and to save space we have not reported them here. We do report results for both functional forms for the other specifications below.

significant. Cooling towers have a surprisingly small estimated effect on costs, although the standard error is large. This probably reflects the difficulty in measuring cooling water control opportunities and responses in more than a very rough way.[34]

The coefficients of the architect-engineer and utility experience variables have the right sign, but are small in magnitude and have standard errors about ten times larger than the coefficient estimates. Using Zimmerman's (1982) functional form for experience does not improve the results.

□ **Heterogeneous technology with equal learning coefficients.** We next estimate the model by allowing the level of costs and the extent of scale economies to differ across technological groups, but constraining learning coefficients to be the same across groups. Architect-engineer experience is measured by experience within each group. We experimented with measuring utility experience in the same way, but generally find that total utility coal experience gives more precise results than does utility experience within individual technological groups. Column 2 reports results with experience measured as the log of cumulative experience, and column 3 reports results with $1/(1 + experience)$.

There are several major differences between these results and the previous "homogeneous technology" results. First, supercritical units exhibit a higher level of costs and larger estimated scale effects than do subcritical units. The null hypothesis that the intercept and scale terms for subcritical and supercritical units are the same can be rejected at the 1% level ($F(2, 347) = 9.680$). The estimated value of construction cost economies of scale for supercritical units is, however, much larger than we expected.

Subcritical units in the 1800 and 2400 psi classes have very similar cost characteristics. The differences in scale coefficients and intercepts for subcritical units with different steam pressure characteristics are not significant at the 5% level ($F(4, 343) = 1.52$). The only anomaly we observe is that for the units in the 2000 psi class, we find no scale economies at all. The imprecision of this result suggests that it may be a consequence of the small sample size—there are only 22 units in this class.

We obtain better results for the new experience variables in the context of this model, which allows for differences in costs across technologies. Both utility and architect-engineer experience have the correct signs and are larger in absolute value than in the homogeneous specification. The effect of architect-engineer experience is more precisely estimated in this specification, although it still has a large standard error. The estimated experience effects appear sensitive to the specification and measurement of technology and associated experience effects. The equation estimated by using Zimmerman's (1982) specification of experience is essentially identical to the one estimated by using the more conventional specification (recall that Zimmerman's functional form implies a positive coefficient for the experience variable). In each equation we can easily reject at the 1% level the null hypothesis that the coal units as a group are drawn from a homogeneous technology ($F(6, 343) = 4.26$ and 4.11).

□ **Heterogeneous technology with variable learning effects.** Finally, columns 4 through 6 report results in which we allow the architect-engineer and utility learning effects to vary across technological groups. The column 6 results include a measure of "industry learning" effects.

Changing the specification in this way has some effect on both the scale coefficients and the experience coefficients, compared with the results reported in columns 2 and 3.

---

[34] There are numerous ways to respond to temperature restrictions on cooling water discharge. Cooling towers of a variety of types may be used, as may cooling ponds, canals, etc. The optimal response and the cost will depend on the site of the plant, ambient water temperatures, variations in these over the year, and other factors. Our cooling tower dummy measures these responses only crudely.

TABLE 3 **Econometric Results (Standard Errors in Parentheses)**

Variable	Equation (1)	Equation (2)	Equation (3)	Equation (4)	Equation (5)	Equation (6)	Equation (7)	Mean	Std. Dev.
RWAGE	.0305	.035	.036	.037	.035	.033	—	10.69	1.13
	(.0127)	(.0126)	(.0126)	(.012)	(.012)	(.0126)			
FIRST	.217	.220	.220	.224	.222	.222	.224	.248	.432
	(.0254)	(.0253)	(.0254)	(.0249)	(.0252)	(.025)	(.025)		
UNCONV	−.044	−.0302	−.033	−.040	−.048	−.049	−.067	.302	.460
	(.0283)	(.0278)	(.0277)	(.028)	(.028)	(.028)	(.030)		
SCRUBBER	.137	.149	.147	.130	.137	.135	.111	.088	.283
	(.045)	(.044)	(.044)	(.0435)	(.044)	(.044)	(.044)		
COOLTWR	.033	.0536	.0530	.0583	.0532	.0541	.017	.302	.460
	(.028)	(.0284)	(.0283)	(.0281)	(.0279)	(.0279)	(.030)		
LSIZE	−.183	—	—	—	—	—	—	5.967	.603
	(.032)								
LSIZE(1800)	—	−.150	−.164	−.220	−.216	−.230	−.225	5.185	.350
		(.0731)	(.0723)	(.0746)	(.0730)	(.0741)	(.075)		
LSIZE(2000)	—	.070	.071	.161	.127	.102	.199	5.606	.380
		(.125)	(.124)	(.138)	(.144)	(.146)	(.140)		
LSIZE(2400)	—	−.135	−.151	−.128	−.149	−.157	−.115	6.015	.430
		(.050)	(.048)	(.050)	(.048)	(.048)	(.050)		
LSIZE(3500)	—	−.452	−.454	−.347	−.366	−.325	−.373	6.534	.320
		(.076)	(.076)	(.084)	(.079)	(.086)	(.085)		
EXPERAE	−.0058	−.0251	.088	—	—	—	—	2.88/.10	.988/.171
	(.0425)	(.019)	(.048)						
EXPERU	−.0010	−.0133	.004	—	—	—	—	1.78/.29	1.083/.302
	(.0176)	(.0175)	(.045)						
EXPERAE(1800)	—	—	—	.0070	.0273	.0312	−.015	1.63/.31	.973/.315
				(.030)	(.0857)	(.0857)	(.031)		
EXPERAE(2000)	—	—	—	−.036	.0635	.0588	−.054	1.25/.43	.991/.354
				(.051)	(.147)	(.147)	(.051)		
EXPERAE(2400)	—	—	—	−.015	.0251	.030	−.024	1.77/.26	.911/.279
				(.024)	(.068)	(.068)	(.024)		

EXPERAE(3500)	—	—	—	−.0954 (.0322)	.266 (.0770)	.233 (.082)	−.088 (.032)	1.32/.37	.806/.310
EXPERU(1800)	—	—	—	.0731 (.030)	−.179 (.068)	−.183 (.068)	.076 (.030)	.89/.54	.824/.353
EXPERU(2000)	—	—	—	−.0825 (.050)	.205 (.180)	.171 (.182)	−.080 (.050)	1.89/.25	1.023/.281
EXPERU(2400)	—	—	—	−.0334 (.021)	.0794 (.0657)	.0723 (.0659)	−.038 (.022)	1.82/.25	.955/.257
EXPERU(3500)	—	—	—	−.0645 (.027)	.239 (.094)	.247 (.095)	−.063 (.028)	2.35/.17	1.067/.228
EXPERI	—	—	—	—	—	.555 (.480)	—	.018	.036
T(1970)	−.422 (.0764)	−.423 (.0710)	−.421 (.0694)	−.400 (.0710)	−.387 (.069)	−.390 (.069)	−.406 (.072)	—	—
S(2000)	—	−1.178 (.755)	−1.260 (.758)	−1.783 (.788)	−2.001 (.917)	−1.920 (.919)	−2.020 (.797)	—	—
S(2400)	—	−.113 (.403)	−.103 (.402)	−.366 (.404)	−.492 (.427)	−.523 (.428)	−.458 (.410)	—	—
S(3500)	—	1.879 (.565)	1.810 (.565)	1.154 (.595)	.735 (.616)	.387 (.685)	1.271 (.607)	—	—
A (mean)	7.0453	6.793	6.878	7.048	7.241	7.346	7.745	—	—
R^2 =	.65	.68	.68	.70	.70	.70	.71		
R^2 (within group) =	.56	.60	.60	.62	.62	.62	.64		

Legend:

Equation (1): Homogeneous technology/log total architect-engineer coal experience/log total utility coal experience.

Equation (2): Heterogeneous technology/equal experience effects/log architect-engineer within technology experience/log total utility coal experience.

Equation (3): Same as equation (2) using (1/1 + *experience*) specification.

Equation (4): Heterogeneous technology/unconstrained experience effects/log architect-engineer within technology experience/log total utility coal experience.

Equation (5): Same as equation (4) using (1/1 + *experience*) specification.

Equation (6): Same as equation (5) adding total industry experience with each technology using (1/1 + *experience*) specification.

Equation (7): Same as equation (4) with dummy variables for BLS regions (regional fixed effects).

All equations are estimated with a separate intercept for each architect-engineer as discussed in the text. We report both a conventional R^2 and a "within group R^2," which measures the proportion of variation within architect-engineer groups explained by the model.

Means and standard deviations reported for the size and experience variables are the values when the relevant variables are active and exclude zero values that appear in the data set when they are inactive. Values for both specifications of the experience variables are reported where appropriate with the log experience specification appearing first.

The absolute magnitude of the coefficient for 1800 psi units increases. We can now reject equality of the intercept and scale coefficients across the subgroups of the subcritical class at the 5% level, but not at the 1% level ($F(4, 337) = 2.45$). Supercritical units continue to exhibit much larger scale economies than do subcritical units, although the size of this coefficient is now smaller and more plausible. We continue to reject homogeneity across subcritical and supercritical technology at the 1% level ($F(4, 345) = 7.15$).

The estimated experience effects generally have the correct signs. Experience effects for supercritical technology are fairly large and are precisely estimated. We can reject at almost the 1% level ($F(2, 341) = 4.26$) the null hypothesis that subcritical and supercritical units have the same experience effects. The one anomaly here is the incorrect sign and statistical significance of utility experience for the 1800 psi class. We suspect that this result may be driven by a small number of very costly 1800 psi units built during the 1970s, and we explore this further below. The estimated industry experience effect has the correct sign, although it is not estimated particularly precisely.

We have performed several sensitivity analyses of these results. First, we estimated the equation by allowing *LSIZE* to enter quadratically as well as linearly. Except for the 2000 psi class, for which we still cannot identify any scale economies, we get negative values for the *LSIZE* coefficients and positive values for the $(LSIZE)^2$ coefficients. This suggests that the elasticity of average cost with respect to size declines as units get larger, a result which is quite plausible. But the $(LSIZE)^2$ terms are never significant either individually or jointly. The estimates of the other coefficients in the model are not affected by introducing the quadratic term. Numerically, within the range of our sample observations for each group, the nonlinear specification traces out essentially the same unit-cost/size relationship as the simpler specification we have been using.

Second, we estimated the cost function with regional fixed effects (i.e., introducing a separate dummy variable for each region but one) to determine whether our estimates are biased by omitted regional characteristics. This precludes estimation of a regional wage effect. These results for the log experience specification are reported in column 7. The coefficient estimates are generally robust to this change in specification. The architect-engineer experience variables are now all negative and estimated a little more precisely. The utility experience variables are essentially unchanged. There is a small but statistically significant overall reduction in the unexplained error. Since we cannot estimate the relationship with both regional dummy variables and a regional wage variable, we prefer to rely on the latter specification, as it allows us to give at least some economic interpretation to regional differences in construction costs.

Next, we explored the specification of experience effects in more detail. We were unable to obtain very precise estimates of the experience effects for the subcritical technology. There are at least two potential explanations of this. There is much more variation in architect-engineer experience for the recent supercritical technology than there is for the much older subcritical technologies during our sample period. (As discussed above, total coal unit experience is used for utilities, so that this should not be a problem for utility experience.) In addition, to the extent that learning effects are important, they may not be constant over long periods of time. Experience effects may simply essentially disappear at some point. In the case of subcritical technology, this may have happened even before our sample period began. Our sample period, however, covers essentially all of the experience with supercritical technology.

Although we cannot identify changes in the experience coefficient that occurred before our sample's starting point, we can test for significant differences in the estimated experience effects within our sample period. To do so, we estimated the model by allowing the experience coefficients for each group to take on different values before and after 1970. For the 2000 psi, 2400 psi, and supercritical groups there is no significant difference in the estimated experience coefficients between the 1960s and 1970s. The anomalous utility experience effect for the 1800 psi group is, however, entirely attributable to units built during the 1970s. During the 1970s most utilities ordered 2400 psi units or supercritical units. The relatively small number of 1800 psi units built in this decade include some of the most expensive units in our sample. Their experience coefficient is probably picking up something unusual about the utilities ordering these units (e.g., they are inefficient), rather than telling us anything about experience effects.

Finally, we estimated the model by allowing for an interaction between architect-engineer experience and utility experience to see whether the combination of architect-engineer and utility experience affected costs. We could find no statistically significant interaction effects.

8. Comparative construction cost relationships

■ **Effects of unit size on construction costs.** We focus our discussion here on the 2400 psi subcritical and the supercritical technologies, which account for about 75% of our sample and represent the primary alternatives from which utilities chose during most of our sample period. The econometric results suggest that there are significant differences in the cost characteristics of these two groups. Figure 1 presents representative estimates of the average real cost per Kw at different scales, evaluated at the means of the independent variables. These curves are labelled 2400 mean and 3500 mean. This allows us to compare the average construction costs of the dominant subcritical technology with those of supercritical technology at different scales. The values presented are average costs in 1980 dollars, *given cost conditions in 1970,* net of standardized interest during construction, for an average architect-engineer. The range of cost estimates reported in each case covers only the scales at which each technology was actually constructed. We have used column 5 in Table 3 above to compute these values, but obtain similar results by using any of the equations reported in columns 4 through 7.

FIGURE 1

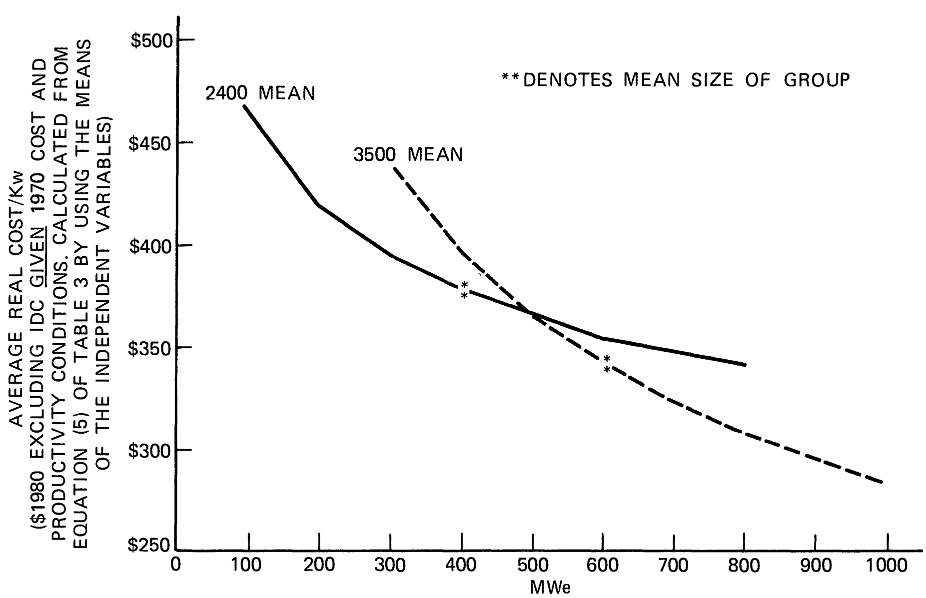

At a scale of 300 Mwe (the smallest size for supercritical units in our sample), supercritical units are over 10% more expensive than subcritical units. The construction cost functions cross at about 500 Mwe, where the average cost of supercritical units falls below that of 2400 psi units. The standard error of the crossover point is about 200 Mwe. The point at which these curves cross is a bit sensitive to the particular specification we choose, but it always occurs between 500 and 600 Mwe. At 700 Mwe supercritical units are about 7% less expensive per Kw than 2400 psi units. Thus, there does not appear to be a simple tradeoff between unit cost and thermal efficiency independent of unit size. Supercritical units (with higher thermal efficiencies) do cost more to build, but not at all scales. Large supercritical units are actually less costly than large subcritical units. There also is no simple "static" tradeoff between unit size and construction cost: full exploitation of economies of scale in construction costs can only be achieved by moving from one technology to another. It would be wrong to think of "static" economies of scale independently of choice of technology.[35]

□ **Experience effects.** We focus on experience effects for supercritical technology, since it is the only technology for which estimated experience effects are always both numerically and statistically significant. For almost all technologies, however, the point estimates suggest experience effects with respect to both architect-engineer and utility experience. Comparing unit cost for a supercritical unit with no architect-engineer or utility experience with unit cost at the means of the experience variables, we find a cost reduction of about 15% associated with both architect-engineer and utility experience. At the maximum experience values, costs are reduced by roughly 20% for each type of experience. There is little quantitative difference between the two functional forms for experience in this range.

□ **Environmental control technology.** On average scrubbers (15%) and cooling towers (6%) have added an estimated 20% to the construction costs of coal units. Given 1980 cost conditions, the presence of a scrubber and a cooling tower increases construction costs by about \$110 per Kw for a 500 Mwe 2400 psi unit, of which about \$80/Kw is for scrubbers and about \$30/Kw is for cooling towers. These costs may be lower than the costs of meeting current air and water pollution regulations and do not represent an estimate of the capital costs of *all* environmental control equipment, much of which we cannot measure directly. But our estimates, especially for scrubbers, are unlikely to be very inaccurate.[36]

□ **Productivity changes/time effects.** One of our most striking and persistent results is the time pattern of costs after controlling for scale effects, technological differences, input price changes, major environmental control technology, and other cross sectional differences in real costs. No matter how one divides the sample or specifies relationships based on equation (1), the resulting pattern of real costs over time, net of input price changes, is

---

[35] See Levin (1977) for a discussion of this issue. The relationship between the costs of 2400 psi subcritical units and supercritical units looks similar to Levin's comparison of "old" and "new" technology for ammonia plants (p. 217). Although we consider only capital costs, while Levin examines total costs, supercritical units are designed to be more fuel efficient than subcritical units (on the order of 5% more fuel efficient) so that the general relationship between the curves in Figure 1 would be similar if we considered total costs of generation (ignoring availability and maintenance problems; see below).

[36] Engineers with whom we have spoken estimate that air and water pollution control equipment has added 20% to 30% to the cost of a typical unit. See Willenbrock and Thomas (1980, p. 798).

FIGURE 2

the same.[37] Real costs per Kw decline during the early and mid 1960s, stabilize in the late 1960s, and then climb during the 1970s to a level that by 1980 is substantially higher than the level in 1960. In Figure 2 we have constructed a cost index for an average 500 Mwe 2400 psi unit with $1980 = 1.0, on the basis of the estimated time effects associated with equation (5) in Table 3. The pattern that emerges is quite typical of all of the specifications reported, and the same pattern emerges when we estimate these relationships by using the undeflated construction cost data.

We do not find this *general* pattern of residual cost changes to be particularly surprising. The conventional wisdom within the industry is that real costs have increased since 1970, owing to increased environmental requirements, longer construction periods, and general declines in construction productivity; this reversed a long historical trend of declining unit construction costs. We were surprised by two things. First, we expected to see the major increases appear later as a result of new plants' coming on line with state-of-the-art environmental control equipment in response to regulations introduced in the 1970s; but costs clearly begin to increase by the late 1960s. This is much too early to attribute the cost increases entirely to unmeasured environmental control costs. Some of the observed cost increase may reflect indirect effects of tightened environmental restrictions, broadly defined to include state and local siting and permitting regulations, that appear as increased construction times and thus higher residual interest during construction costs. But the cost increases appear at least partially to reflect a general decline in construction productivity in this industry, independent of the costs of adding new environmental control equipment.[38]

The second surprise is the magnitude of the estimated cost increases from trough to peak. Commentators on the industry often make cost comparisons by using nominal

[37] We have estimated this relationship separately for each technological group and for different time periods. The pattern of time effects is very robust.

[38] Declining construction productivity has been of concern in the utility industry as well as in other industries. See Business Roundtable (1983).

dollars and without correcting for changes in interest rates over time, differences in generating technology, or scrubber and cooling tower costs.[39] We had anticipated that a large part of the conventional wisdom regarding increased construction costs represented a combination of "inflation illusion" and "technology illusion." After accounting for input price changes and the technological characteristics of the units, we expected to find a relatively small increase in residual costs that reflected unmeasured environmental control equipment, construction delays, and perhaps some general decline in construction productivity. Instead, the increase in residual cost from the late 1960s to the late 1970s is about 80%. Adding the costs of scrubbers and cooling towers brings it to about 100%. Perhaps these cost patterns can be tied in some indirect way to responses to changing environmental regulations which we cannot measure directly, but we suspect that the answer lies in part in more general problems of productivity in construction.

9. What happened to supercritical technology?

■ Before concluding, we want to identify a puzzle that emerges from our empirical results. As we discussed earlier, the pattern of increasing steam temperatures and pressures that new generating units had exhibited through time stopped in the mid-1970s and now appears to have been reversed. Supercritical technology, which seemed so promising in the 1960s and which accounted for 60% of new capacity by the mid-1970s, has been almost abandoned today. In addition, the average size of new units has gradually declined since the mid-1970s.

These patterns appear to be inconsistent with the estimated construction cost relationships. Larger units continue to be less costly to build than smaller units. Furthermore, large-scale supercritical units, which are theoretically more fuel efficient than subcritical units, are, for a reasonably experienced architect-engineer and utility, less expensive to build than comparable subcritical units. For utilities building relatively large units, supercritical technology would appear to be the preferred choice on both capital cost and fuel cost grounds. Even if supercritical units were slightly more expensive than subcritical units, increased fuel efficiency could justify choosing supercritical technology over subcritical technology.

Although we cannot yet provide a definitive explanation of this puzzle, we want to suggest what appear to us to be the two possible reasons for these changes and to present some evidence that is consistent with one of them. A more complete analysis is the subject of ongoing research. Two factors may have motivated these changes in utility investment behavior. First, electricity demand expectations have declined considerably during the 1970s. In the early 1970s utilities' planning assumed average demand growth of about 8% per year. By the early 1980s, this had dropped to less than 3% per year.[40] As a result, building ahead of demand to achieve the construction cost economies of scale and the increased theoretical thermal efficiency of large supercritical units has become substantially less attractive to utilities.[41]

---

[39] Recall that the nominal cost figures include accumulated interest charges. In the late 1970s high nominal interest rates led to reported construction cost figures that had a relatively high proportion of interest-during-construction charges.

[40] See Table 5, "Average Annual Rates of Growth in Summer Peak Load Forecasted by the Electric Reliability Councils, 1967–81," in "Nuclear Plant Cancellations: Causes, Costs and Consequences," U.S. Department of Energy, DOE/EIA-0392, (April 1983), p. 16.

[41] See Snow (1975) for a review of the literature on investment cost minimization in a growing system with scale economies.

Demand-side considerations no doubt account for some of the movement away from large units and the simultaneous shift toward 2400 psi units, which are more economical at small scale. But we do not believe that demand-side considerations alone can explain recent patterns of unit additions. Although the average size of units has declined, and very large units (greater than 750 Mwe) are rare, there were still 26 units (24 of which were subcritical units) greater than 500 Mwe installed in 1981 and 1982. The reduction in average unit size is largely, but not completely, associated with the movement away from supercritical units which dominated the larger units as a whole during the 1970s. We believe that there are probably supply-side considerations other than construction costs that have led utilities to eschew supercritical units in particular and very large units more generally. These are the relatively poor reliability and high maintenance costs of supercritical units, which were not expected when initial commitments to supercritical technology were made.[42]

Gordon's (1982) interviews with plant operators revealed that costly maintenance problems were experienced by operators of state-of-the-art supercritical units. High maintenance costs are normally associated with poor plant reliability as well. Poor plant reliability increases replacement power costs. Thus, although large supercritical units are, on average, somewhat less expensive than similar subcritical units, poor plant performance may increase the *total* costs of generation from these units and make them unattractive investments.

Reliability data that we have collected and analyzed in connection with a related study (Joskow and Schmalensee, 1985) are consistent with Gordon's interview responses. In Table 4 we report estimates of expected average equivalent availability for subcritical and supercritical units by size, for a sample of over 250 coal units. These estimates incorporate adjustments for differences in unit age, cycling characteristics, and coal characteristics.[43] On average, supercritical units have much lower availabilities than do

TABLE 4 **Unit Equivalent Availability by Type and Size of Units***

Size	Subcritical	Supercritical
100	82.8	—
200	80.3	—
300	77.7	70.5
400	75.2	67.6
500	72.7	65.0
600	70.2	62.9
700	67.6	61.1
800	65.1	59.7
900	62.6	58.6
1000	—	57.9

* Estimates based on equivalent availability data for 276 units over the period 1969 to 1980, adjusted for unit age, initial operating date, and coal characteristics.

[42] Perl (1982) makes the first effort to integrate these considerations in an analysis of total life-cycle costs of coal-fired generation.

[43] These estimates are based on an analysis of a time-series cross section of operating performance for a sample of coal units. The analysis involves GLS estimation of equations relating equivalent availability to a variety of unit characteristics.

subcritical units. Especially poor performance is exhibited by the larger supercritical units, although large units generally have lower availabilities than small units.

These data suggest that one likely reason for the demise of supercritical technology, as well as the general movement away from very large units, is that the construction cost advantages at large scale and the theoretical thermal efficiency advantages of supercritical units are dominated by the increased maintenance costs and replacement power costs required to compensate for reduced effective capacity. Large units may be cheaper to build than smaller units; and supercritical units may yield higher theoretical thermal efficiencies and be less expensive to construct than are subcritical units. But if the units do not work very well, the lower construction costs are not a good indicator of the effective *total* cost of electricity.

10. Conclusions

■ Our analysis of the costs of building coal-burning generating units during the period 1960 to 1980 leads to the following conclusions.

(1) Contrary to some previous econometric work and popular "conventional wisdom," there are significant economies of scale associated with generating unit construction costs. The extent of economies of scale in construction is technology-specific, however. The assumption of homogeneous technology with respect to the level and scale effects associated with construction costs is rejected. The econometric results are broadly consistent with the assumptions made in engineering cost calculations.

(2) We find evidence that both architect-engineer experience and utility experience lead to lower construction costs. The estimated experience effects are numerically and statistically more significant for supercritical technology than for subcritical technology. The failure to distinguish between different technologies leads to an underestimate of these learning effects. In conjunction with economies of scale in construction costs, the presence of learning effects, especially utility-specific learning effects, suggests that there are likely to be cost advantages for larger firms that may not be achieved, given the current fragmented nature of the electric utility industry.

(3) The major identifiable construction responses to air and water pollution regulations (scrubbers and cooling towers) have added at least 20% to the real construction costs of coal units. These costs are only a small fraction of the total increase in real construction costs observed for the 1970s, however.

(4) After accounting for the primary factors that are generally thought to influence static construction costs and adjusting for changes in input prices, we find a large pervasive increase in residual real costs since the late 1960s. Other things equal, real costs have increased by 80% since the late 1960s. These increases appear to reflect the costs of responding to environmental regulation not otherwise accounted for in the specification, increased construction times, and declining construction productivity. The precise reasons for such large real-cost increases remain a puzzle and a subject for future research.

(5) The apparent "technological frontier" in the 1960s—supercritical technology—seems to have been abandoned by the early 1980s. The reason for this is not unexpectedly high construction costs *per se,* however. The demise of supercritical technology appears to reflect much smaller demands for new capacity, which make large units generally less economical, and poor operating performance, with its associated high maintenance and replacement power costs. Poor operating performance appears to have reduced the economic attractiveness of large units generally.

Appendix: the data

■ The primary data base includes all 440 units listed in the Department of Energy's *Generating Unit Reference File* (*GURF*) which began operating during the years 1960 to 1980, were designed to burn coal, and actually did burn coal in their first year of operation. In addition to the primary data base, we have data from a census that we conducted of all coal units built between 1950 to 1959 inclusive. For these units we have information on unit size, operating year, utility, turbine throttle pressure, and architect-engineer. These data were used along with the information in the primary data base to construct the experience variables for the 1960–1980 units and to construct Tables 1 and 2. Finally, we have identified all coal units which began operating in 1981 and 1982, and their sizes and steam pressure characteristics, to update Tables 1 and 2. Construction cost data and other characteristics of these last units were not available when this article was written. Of the 440 units in the primary data base, we have information on all variables in the estimating relationships for 411. The missing information for the others is typically the construction cost.

We used a variety of sources to identify the units for the primary and secondary data bases. These included *GURF*, the *Inventory of Power Plants* (various years), the Federal Energy Regulatory Commission's (formerly the Federal Power Commission) *Steam-Electric Plant Construction Costs and Annual Production Expenses* (annual, 1949 to 1980—FPC statistics), annual surveys in *Power*, and the NUS Corporation's *Commercial Coal Power Plants* (*NUS*).

Construction cost data are from *Steam-Electric Plant Construction Cost and Annual Production Expenses* (annual, 1959 to 1980). This publication reports the gross expenditures on plant and equipment (including structures and land) by generating plant. To obtain construction costs of each unit we identified the year that a particular unit began commercial operation. The expenditures on a unit enter the gross expenditure account for the plant in the year that it begins commercial operation. By subtracting the construction expenditures listed for the previous year from those reported for the year a plant begins commercial operation, we obtain an estimate of the "as spent" nominal construction costs for that unit. A potential source of error in this estimate would occur if the existing units incurred large capitalized maintenance expenditures in the same year that a new unit comes on line. In some cases, two units of a plant come on line in the same year, and there is no way to identify separate construction costs for each. When this happens the units are almost always identical. We assign each unit the same construction cost. In this case, if one of the units is the first unit on the site, we do not give it a value of one for the first unit dummy.

We adjust the nominal construction cost data to reflect changes in input prices and interest rates, and express the deflated costs in $1980, excluding standardized interest during construction charges. For each unit we start with the reported "as spent" nominal construction cost and use the following formula to adjust the reported figure to net out interest charges and to adjust the reported expenditures to reflect prices prevailing during the initial year the unit operated:

$$\begin{array}{l}\textit{total cost in initial year in}\\ \textit{constant dollars (net of}\\ \textit{interest during construction)}\end{array} = \frac{\textit{as spent nominal cost}}{\sum_{t=1}^{5} S_t \cdot [\prod_{i=1}^{t} (1 + p(i)) \cdot \prod_{j=t}^{5} (1 + r(j))]}, \qquad \text{(A1)}$$

where

S_t = the share of actual construction expenses in year t, taken from a typical cash flow curve as described in *Power Plant Capital Costs, Current Trends and Sensitivity to Economic Parameters,* U.S. Atomic Energy Commission, WASH-1345, October 1974, Figure 5. This gives annual cash flows for a five-year construction period.

$P(i)$ = the percentage change in input prices in year i, taken from the *Handy-Whitman Public Utility Construction Cost Index for All Steam Plants* (by region).

$r(j)$ = the average allowance for funds used during construction rate from the Department of Energy's *Statistics of Privately-Owned Utilities in the United States* (various years—earlier editions produced by the Federal Power Commission).

Nameplate generating capacity for each unit is based on the capacity reported in the *FPC Statistics* for the year the unit was placed in service. For the few units not reported there, we used the nameplate capacities reported in *GURF*. Steam pressure and indoor/outdoor construction characteristics were obtained primarily from the *FPC Statistics*. This source was supplemented with information drawn from *Power, NUS,* and *Electrical World*. The cooling tower information was obtained from *GURF*.

The regional wage data are the reported average union wage plus employer benefit contributions for helpers and journeymen in the building trades. This information was obtained from the Bureau of Labor Statistics, *Union Wages and Hours: Building Trades,* July 1, 1976, Bulletin 1972, Table 12.

Information on the architect-engineer for each unit was obtained from a variety of sources including *NUS, Power,* the Perl data base (see below), and interviews of electric utilities. The scrubber information was obtained from the *EPA Utility FGD Survey, October-December 1979* (January 1980), *GURF,* and *NUS.*

The data collection effort, especially the tedious task of collecting the construction cost information, was eased substantially by Lew Perl who made his data for 245 coal units available to us. We checked a subset of the data in his sample both for accuracy and to ensure that we were using a symmetrical computation procedure for the other 200 units. We use a somewhat different computation procedure for developing the $1980 cost figures; our cash flow pattern is slightly different, and we use regional input price indexes.

References

ALCHIAN, A. "Reliability of Progress Curves in Airframe Production." *Econometrica,* Vol. 31 (October 1963).

THE BUSINESS ROUNDTABLE. "Summary Report of the Construction Industry Cost Effectiveness Project." New York: January 1983.

BUSHE, D. "An Empirical Analysis of Production and Technology Using Heterogeneous Capital: Thermal Electric Power Generation." Unpublished Ph.D. Dissertation, New York University, 1981.

COOTNER, P.H. AND LOF, G.O.G. *Water Demand for Steam Electric Generation.* Washington, D.C.: Resources for the Future, 1965.

COWING, T.G. "Technical Change and Scale Economies in an Engineering Production Function: The Case of Steam Electric Power." *Journal of Industrial Economics,* Vol. 23 (December 1974), pp. 135–152.

——— AND SMITH, V.K. "The Estimation of a Production Technology: A Survey of Econometric Analysis of Steam Generation." *Land Economics,* Vol. 54 (May 1978), pp. 156–186.

DOW CHEMICAL COMPANY. "Energy Industrial Center Study." June 1975.

EDISON ELECTRIC INSTITUTE. *Statistical Yearbook of the Electric Utility Industry/1982.* Washington, D.C.: October 1983.

ELECTRIC POWER RESEARCH INSTITUTE. *Technical Assessment Guide,* EPRI P-2410-SR. May 1982.

FRENCH, R.X. AND HADDAD, S.Z. "The Economics of Reliability and Scale in Generating Unit Size Selection," Paper GEC-P196. Sargent & Lundy, Inc., March 1980.

GOLD, B. "On Size, Scale and Returns: A Survey." *Journal of Economic Literature,* Vol. 19 (March 1981), pp. 5–33.

GORDON, R.J. "The Productivity Slowdown in the Steam-Electric Generating Industry." Unpublished Manuscript, October 1982.

HAUSMAN, J.A. "Specification Tests in Econometrics." *Econometrica,* Vol. 46 (November 1978).

——— AND TAYLOR, W.E. "Panel Data and Unobservable Individual Effects." *Econometrica,* Vol. 49 (November 1981).

HIRSCH, W. "Manufacturing Progress Functions." *Review of Economics and Statistics,* Vol. 34 (May 1952).

JOSKOW, P.L. "Vertical Integration and Long-Term Contracts: The Case of Coal-Burning Power Plants." *Journal of Law, Economics, and Organization,* Vol. 1 (1985).

——— AND SCHMALENSEE, R. *Markets for Power: An Analysis of Electric Utility Deregulation.* Cambridge: MIT Press, 1983.

——— AND ———. "The Performance of Coal-Burning Generating Units in the U.S.: 1960–1980." Mimeo, May 1985.

KOMANOFF, C. *Power Plant Cost Escalation.* New York: Komanoff Energy Associates, 1981.

LEVIN, R. "Technical Change and Optimal Scale: Some Evidence and Implications." *Southern Economic Journal,* Vol. 44 (1977), pp. 208–221.

LIEBERMAN, M.B. "The Learning Curve and Pricing in the Chemical Processing Industries." *Rand Journal of Economics,* Vol. 15 (Summer 1984), pp. 213–228.

LING, S. *Economies of Scale in the Steam Electric Power Generating Industry.* Amsterdam: North-Holland, 1964.

Mark's Standard Handbook for Mechanical Engineers, 8th ed.

MOOZ, W.E. "Cost Analysis of Light Water Reactor Power Plants." Santa Monica: The Rand Corporation, June 1978.

NERLOVE, M. "A Note on Error Components Models." *Econometrica,* Vol. 39 (March 1971).

PERL, L. "The Current Economics of Electric Generation from Coal in the U.S. and Western Europe." Unpublished Manuscript, National Economic Research Associates, October 1982.

OAK RIDGE NATIONAL LABORATORY. "ORCOST—A Computer Code for Summary Cost Estimates of Steam Electric Power Plants." ORNL-TM-3743. September 1972.

SCHMALENSEE, R. AND JOSKOW, P.L. "Estimated Parameters as Independent Variables—with an Application to the Costs of Electric Generating Units." Sloan School of Management Working Paper, February 1985.

SEARLE, A.D. AND GOODY, C.S. "Productivity Increases in Selected Wartime Shipbuilding Programs." *Monthly Labor Review* (1945), pp. 1132–1142.

SNOW, M.S. "Investment Cost Minimization for Communications Satellite Capacity: Refinement and Applications of the Chenery-Manne-Srinivasin Model." *Bell Journal of Economics,* Vol. 6 (Autumn 1975), pp. 621–643.

STEWART, J.F. "Plant Size, Plant Factor, and the Shape of the Average Cost Function in Electric Power Generation: A Nonhomogeneous Capital Approach." *Bell Journal of Economics,* Vol. 10 (Autumn 1979), pp. 549–565.

U.S. DEPARTMENT OF ENERGY, "Delays and Cancellations of Coal-Fired Generating Capacity," DOE/EIA-0406. July 1983.

WILLENBROCK, J. AND THOMAS, H.R. *Planning, Engineering, and Construction of Electric Power Generating Facilities.* New York: John Wiley, 1980.

WILLS, H.R. "Estimation of a Vintage Capital Model for Electricity Generating." *Review of Economic Studies,* Vol. 45 (October 1978).

ZIMMERMAN, M. "Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power." *Bell Journal of Economics,* Vol. 13 (Autumn 1982), pp. 297–310.