

Historical Costs of Coal-Fired Electricity and Implications for the Future

James McNerney^{*,a,b}, Jessika E. Trancik^{*,b,d}, J. Doyne Farmer^{b,c}

^aDepartment of Physics, Boston University, Boston, MA 02215, USA

^bSanta Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

^cLUISS Guido Carli, Viale Pola 12, 00198, Roma, Italy

^dEngineering Systems Division, Massachusetts Institute of Technology, Cambridge, MA 02139-4307, USA

Abstract

We study the costs of coal-fired electricity in the United States between 1882 and 2006 by decomposing it in terms of the price of coal, transportation costs, energy density, thermal efficiency, plant construction cost, interest rate, and capacity factor. The dominant determinants of costs at present are the price of coal and plant construction cost. The price of coal appears to fluctuate more or less randomly while the construction cost follows long-term trends, decreasing from 1902 - 1970, increasing from 1970 - 1990, and leveling off or decreasing a little since then. This leads us to forecast that even without carbon capture and storage, and even under an optimistic scenario in which construction costs resume their previously decreasing trending behavior, the cost of coal-based electricity will drop for a while but eventually be determined by the price of coal, which varies stochastically but shows no long term decreasing trends. Our analysis emphasizes the importance of using long time series and comparing different electricity generation technologies using total costs, rather than costs of single components like capital.

Contents

1	Introduction	1	5.4	Capital and O&M costs	12
2	Historical data	2	5.5	Fuel costs	13
2.1	Data sources	2	5.6	Extrapolation summary	14
2.2	Decomposition formula	2	5.7	Extrapolation results	14
2.3	Fuel cost	3	6	Conclusions	14
2.3.1	Coal price	3	7	Acknowledgments	15
2.3.2	Transportation costs	3	A	Change decomposition	15
2.3.3	Coal energy density	3	1.	Introduction	
2.3.4	Thermal efficiency	4			
2.3.5	Fuel cost	5			
2.4	Operation and maintenance cost	5			
2.5	Capital cost	5			
2.5.1	Specific construction cost	5			
2.5.2	Capacity factor	6			
2.5.3	Interest rate	6			
2.5.4	Capital cost	7			
2.6	Total cost	7			
3	Analysis of cost trends	8			
4	Analysis of cost variation	9			
5	Extrapolation	10			
5.1	Commodities vs. pure technologies	11			
5.2	Modeling a commodity	11			
5.3	Modeling a technology	12			

*Corresponding author

econometrics, where generation costs, or more typically plant costs, are broken down using a regression model [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13].

We also focus on the long term evolution of quantities averaged across plants in the United States, going back to the earliest coal-fired power plant in 1882 through 2006, rather than cross-sections or panels of plants. Thus, our data set sacrifices cross-sectional richness to examine a longer time span (over a century). This is important for characterizing the factors driving the long term evolution of costs.

The data suggests a qualitative difference between the behavior of fuel and capital costs, the two most significant contributors to total cost. Coal prices, which became the most important determinants of fuel costs when average thermal efficiencies ceased improving in the U.S. during the 1960s, have fluctuated and shown no overall trend up or down. This fluctuation and lack of trend are consistent with the facts that coal is a traded commodity and therefore it should not be possible to make easy arbitrage profits by trading it. According to standard results in the theory of finance, this implies that it should follow a random walk. In contrast, plant construction costs, the most important determinant of capital costs, followed a decreasing trajectory until 1970, consistent with what one expects from a technology. This pattern continued until 1970 when construction costs dramatically reversed direction, at least in part due to pollution controls.

To probe the future of coal-fired electricity, we study its costs in a reasonable best case scenario. We study a low-cost scenario for future generation costs involving a return of plant cost improvement to its pre-1970 rate. Under such a scenario, we estimate that capital costs would drop 67% over 50 years, reducing total generation costs by 40%. At the same time, fuel costs would rise in their relative contribution to the total, leading to higher uncertainty in total generation costs, and creating a floor below which the generation cost would be unlikely to drop.

Our analysis makes several methodological advances. It calculates total costs of generation, rather than the costs of individual components like capital, and over a long (~ 100 year) time span. We repeat and refine an approach used by Nemet [14] to decompose changes in cost. The refinement eliminates artificial residuals arising from the cross-effects of variables influencing the cost. Finally, we note that physically accurate models like the one presented here and in Nemet may sometimes allow (data permitting) for more reliable decompositions of cost than regression models.

The paper is organized as follows: in Section 2, we present the historical record of costs for coal-fired electricity, as represented by time series for a number of important variables influencing generation cost. The generation cost consists of three main *components* – fuel, capital, and operation and maintenance. The first two are further decomposed; the fuel component into the coal price, transportation cost, coal energy density, and thermal efficiency; the

capital component into plant construction costs, capacity factor, and interest rate. In Section 3, we determine the variables contributing most to the historical changes in generation costs, focusing on their long term trends rather than their short term variation. In Section 4, we examine the effect their short term variation has on the variation of total generation costs. In Section 5, we use the experience gained from analyzing the data to construct two possible extrapolations of coal-fired generation costs.

2. Historical data

2.1. Data sources

Sources for data are given in the captions of the relevant figures. The data comes from the U.S. Energy Information Administration, Census Bureau, Bureau of Mines, Federal Energy Regulatory Commission (formerly the Federal Power Commission), Federal Reserve, the Edison Electric Institute, and the Platts UDI world electric power plants database, with some minor contributions from a few other databases and technical reports. Methods for handling specific data issues are described in the Appendix and in the Supporting Information. Assumptions made to fill in for missing data are noted in the relevant figures and text.

All data presented here are for the United States between 1882 (the approximate beginning of the electricity utility industry in the U.S.) and 2006. All data are estimates for averages over U.S. coal utility plants. All prices and costs are presented in real 2006 currency, deflated using the GDP deflator. The data is made available at www.santafe.edu/files/coal_electricity_data.

2.2. Decomposition formula

We seek to build up the total generation costs from the following variables, for which data is available:

- OM = operation and maintenance cost (¢/kWh)
- COAL = price of coal ($\text{\$/ton}$)
- TRANS = price of transporting coal to plant ($\text{\$/ton}$)
- ρ = energy density of coal (Btu/lb)
- η = plant efficiency
- SC = specific construction cost ($\text{\$/kW}$)
- r = nominal interest rate
- CF = capacity factor

Specific construction cost (SC) here means the construction cost of a plant per kilowatt of capacity. Each variable is given a subscript t to denote its value in year t . The total generation cost (TC) in year t is

$$\begin{aligned} \text{TC}_t &= \text{OM}_t + \text{FUEL}_t + \text{CAP}_t \quad (1) \\ &= \text{OM}_t + \frac{\text{COAL}_t + \text{TRANS}_t}{\rho_t \eta_t} + \frac{\text{SC}_t \times \text{CRF}(r_t, n)}{\text{CF}_t \times 8760 \text{ hrs/year}} \end{aligned}$$

The three main terms are the three major cost components – operation and maintenance, fuel, and capital. The fuel component accounts for the cost of coal and its delivery to the plant, the amount of energy contained in coal, and the efficiency of the plant in converting stored chemical energy to electricity. The capital component levelizes the construction cost of the plant using the capital recovery factor, $CRF(r_t, n)$, defined as

$$CRF(r_t, n) = \frac{r_t(1 + r_t)^n}{(1 + r_t)^n - 1}, \quad (2)$$

where n is the plant lifetime in years. The capital recovery factor is the fraction of a loan that must be paid back annually, assuming a stream of equal payments over n years and an annual interest rate r_t . For a plant of capacity K , the capital component is the annuity payment on money borrowed for construction, $K \times SC_t \times CRF(r_t, n)$, divided by the yearly electricity production of the plant, $K \times CF_t \times 8760$ hrs/year. Note that K cancels out. Thus, the capital component is the annuity payment on borrowed plant construction funds *per* kilowatt-hour of annual electricity production by the plant. We follow convention and levelize over an assumed plant lifetime of $n = 30$ years.

Note that the total cost for year t uses the capital cost of plants built that year, while the actual fleet of plants existing in year t also includes plants built in previous years. Since we do not have data on the retirement dates of plants, we use the cost of the plants built in year t to estimate the cost of electricity generated in that year. Average capital costs of the whole fleet existing in a particular year would have the same basic evolution as that of new plants alone, but tend to lag the latter, since they would include capital costs from previous years.

In Sections 2.3 – 2.5, we present and discuss the time-series data for each of these variables separately. In Section 2.6 we combine the variables to compute the total cost, and examine the effect that individual variables had on it. The ensuing sections are an extensive discussion of the data; readers more interested in final analyses may wish to skip to Section 3.

2.3. Fuel cost

2.3.1. Coal price

Real coal prices (COAL) have varied over the past 130 years (Fig. 1.) The largest change occurred between 1973 and 1974. A government study [15] found this was due to the OPEC oil embargo started in December 1973, which raised prices for substitutes coal and natural gas, and anticipation of a strike by the United Mine Workers union, which later occurred at the end of 1974. Wage increases starting in 1970 also played a somewhat less important role.

Coal prices at the mine were derived from price time-series for individual varieties of coal. For the period 1882-1956, the coal price is a production-weighted average over anthracite and bituminous varieties, and for 1957-2006




Figure 1: The price of coal at the mine, the price of transporting it to the plant, and their sum. The price of coal has fluctuated with no clear trend up or down. Transportation contributed a significant while decreasing expense to fuel costs. Source: [16, 17, 18, 19, 20]

over anthracite, bituminous, subbituminous, and lignite. No data could be found to weight prices by the quantities in which they were consumed by plants, and therefore we relied on the quantities in which they were produced. The resulting average price series closely resembles that of bituminous coal by itself.

2.3.2. Transportation costs

Transportation costs (TRANS) were derived mainly from the cost of coal delivered to power plants and the price of coal at the mine, though some direct data also exists. Transportation costs have added a significant though decreasing amount to the cost of coal delivered to power plants (Fig. 1). Transportation costs before 1940 were on par with the price of coal at the mine; since 1940, they have dropped to 20 – 40% of the delivered cost on average, though there is considerable variation from plant to plant. The introduction of unit trains in the 1950s may account for some of the decrease in transportation costs. A unit train carries a single commodity from one origin to one destination, shortening travel times and eliminating the confusion of separating cars headed for different destinations. About 50% of coal shipped in the United States (90% of which is used by utilities) is carried by unit trains [21]. Other means of transporting coal are barges, collier ships, trucks, and conveyor belts in cases where plants are built next to mines.

Many utilities responded to environmental regulations by switching to higher priced, low-sulfur coals [10]. Switching to low sulfur coals also extended transportation distances, which may explain the increase in transportation costs seen around 1970.

2.3.3. Coal energy density

Since 1960, the average energy density of coal (ρ) has dropped steadily (Fig. 2 inset.) The lower the energy density, the more coal needed by plants to consume equal amounts of primary energy. Thus, the effect of lower en-

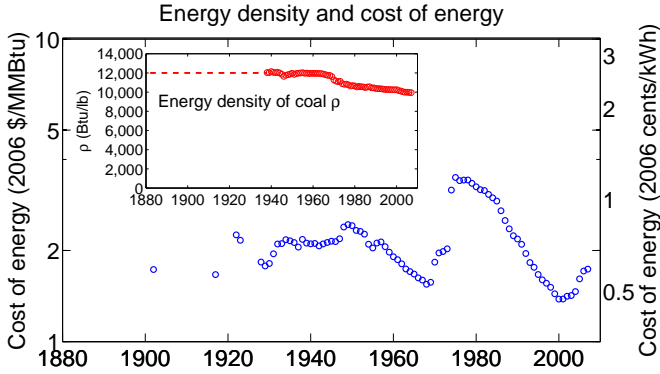


Figure 2: The cost of coal-stored energy, $(\text{COAL}_t + \text{TRANS}_t)/\rho_t$. The left axes show cost in dollars per million BTU, the right axes in cents per kWh. The inset shows the energy density of coal in Btu per pound. The dashed line indicates an assumed value of 12,000 Btu per pound. Coal energy densities have come down about 16% in the last 40 years. Source: [22, 17]

ergy densities has been to increase fuel costs through increased purchase and transportation costs.

Changes in the energy density reflect changes in the overall mixture of coal species used by the industry, as well as variation of energy density within species. One cause of the decrease in energy density may be the increased use of subbituminous coal, which burns more cleanly than previously used bituminous due to a lower sulfur content, but also contains less energy per pound. No data for energy density could be found before 1938, and a value of 12,000 Btu per pound was assumed (the 1938 value).

2.3.4. Thermal efficiency

The thermal efficiency of a fossil-steam plant (η) is the fraction of stored chemical energy in fuel that is converted into electrical energy (Fig. 3). It accounts for every effect which causes energy losses between the input of fuel and the bus-bar of the plant, the point where electricity enters the electric grid. These effects include incomplete burning of fuel, radiation and conduction losses, stack losses, excess entropy produced in the turbine, friction and wind resistance in the turbine, electrical resistance, power to run the boiler feed pump and other plant equipment, and the thermodynamic limits of the heat cycle used.

Over the first 80 years, the average efficiency of coal plants increased by more than a factor of 10. The earliest plants obtained efficiencies below 3%. Improving efficiencies meant that less coal was required to produce equal amounts of electricity, lowering fuel costs. For the last 50 years, the average efficiency of coal-fired plants has stayed approximately level at 32–34%, although individual plants obtained efficiencies as high as 40% during the 1960s [23].

The factors influencing efficiency of plants fall into two major categories: thermodynamic factors, which pertain to the specific heat cycle employed, and operational factors, which reflect the mechanical and electrical efficiencies




Figure 3: Average efficiency of coal plants and cost of the fuel component, $(\text{COAL}_t + \text{TRANS}_t)/(\rho_t \eta_t)$. Efficiency increases drove decreases in fuel costs until 1960. Since then efficiency has remained stable around 32 – 34%. Source: [24, 17]

of individual stages and components. The thermodynamic limit efficiency is determined by the details of the steam cycle used – the maximum and minimum temperatures, the maximum and minimum pressures, and the exact type of cycle followed (e.g. the number of reheat stages used.) Current plant designs typically have thermodynamic limit efficiencies around 46%.¹ The operational efficiencies of the boiler, turbine, generator, and other components further reduce the efficiency achievable in practice (Table 1.) The net effect of the operational efficiencies is to reduce total efficiency by a factor of about 0.7.

Historical increases in efficiency came from improvements in both the thermodynamic and operational categories. Higher steam temperatures, higher steam pressures, and changes to the heat cycle relaxed the thermodynamic constraint, while better designed parts reduced operational losses. However, operational factors are now working at high efficiencies (Table 1), while changing the parameters of the heat cycle faces highly diminishing returns for efficiency [23]. This latter problem reduces incentives to develop economical materials that could withstand higher pressures and temperatures. The overall result has been a standstill in average efficiency for the last 50 years. Given the present technological options, the current aver-

¹We assume a Rankine cycle with superheating and one reheat cycle operating at 1000 °F, 2400 psi.

Table 1: Operational sources of energy loss. Source: [23]

Source of energy loss	Efficiency
Stack losses, radiation & conduction from boiler	87%
Excess entropy produced in turbine	92%
Windage, friction, & elec. resistance	95%
Boiler feed pump power requirement	95%
Auxiliary power requirements	97%
Total (product)	70%




Figure 4: The cost of operation and maintenance of coal plants. A value equal to 21% of the overall fuel cost for a given year was used to fill in missing data. Source: [22, 25, 26, 27, 28, 29]

age efficiency presumably represents an optimum after balancing the higher fuel cost of a less efficient plant against the higher construction costs of a more efficient plant.

2.3.5. Fuel cost

Over the entire time period (1882-2006) the factors most responsible for changes in the fuel cost are the price of coal and the efficiency. Changes in energy density and transportation costs had relatively minor effects. As can be seen in Fig. 3, the net effect of varying coal price, energy density, transportation cost, and efficiency on the fuel component of electricity costs was a long term decrease in the cost of the fuel component until about 1970. The decrease was mainly due to improving efficiency. After 1970, coal prices increased dramatically during a time when efficiency was flat, increasing overall fuel costs 70% between 1970 and 1974.

2.4. Operation and maintenance cost

The O&M cost (OM) is the least significant of the three cost components (fuel, O&M, capital), representing about 5-15% of total generation costs during the last century (Fig. 11). Although it is less significant than fuel or capital, we attempted to reconstruct a time-series for O&M.

Reliable historical data for O&M costs was difficult to acquire. Sources were frequently in conflict with each other due to differences in definition of O&M costs. We use data from the Energy Information Administration (EIA) because it uses a consistent definition over the longest time period. To fill in missing years, we took O&M costs to be 21% of the value of the overall fuel cost, based on the empirical observation that O&M costs were consistently about 21% of overall fuel costs between 1938 and 1985 in all years for which we had data.² Later O&M costs break

²O&M's share was always between 19 and 24%, with an average of 21%, during three periods for which data was available: 1938-1947, 1956-1963, and 1978-1985.




Figure 5: Specific capital cost of coal plants. Plant costs generally decreased until 1970, when pollution regulation, increased construction times, and possibly other factors caused a rise. Source: [33, 34, 35]

from this pattern. This assumption seemed the best choice to avoid introducing discontinuities or other artificial behaviors in the total costs. See the Supporting Information for further details.

Based on our reconstruction, O&M costs decreased until 1970 (Fig. 4). For pulverized coal plants, historical declines “are attributed mainly to the introduction of single-boiler designs, automatic controls, and improved instrumentation” [30, 31]. From the data available, we infer that an increase occurred between 1960 and 1980. This is supported by observations that pollution controls introduced to plants during this time raised O&M costs [32, 11].

2.5. Capital cost

2.5.1. Specific construction cost

Construction cost data were obtained from *Historical Plant Cost and Annual Production Expenses 1982* [33], published by the Energy Information Administration. A potential source of bias with this data is that plant costs were given as accounts which accumulate a utility’s nominal expenditure on a given plant. These accounts therefore include costs of additional units since its construction, and sum together nominal expenditures from different years.

The specific construction cost (SC) is the cost of building a plant per kW of capacity installed. Specific construction costs decreased from the beginning of the industry until about 1970, then doubled between 1970 and 1987 (Fig. 5.) They appear to have been roughly flat since then, though we lack data for the period 1988-1999.

Construction costs have been a focus of interest after 1970 because of this uptick. Several contributing factors have been identified. One of them is pollution controls. Starting with the Clean Air Act in 1970, several laws were passed in the US, requiring plants to add pollution controls to reduce the level of NO₂, SO_x, and particulates from flue gas emissions [11, 10]. The Clean Air Act of 1970 was followed by a number of similar acts between 1970 and 2003. While some utilities initially responded to pollution controls by switching to low-sulfur coals, others responded by installing de-sulfurization equipment known as

sulfur scrubbers. Eventually all new plants were required to install scrubbers [10]. This new equipment raised O&M costs and decreased efficiency slightly, but mainly raised construction costs.

Besides raising plant costs directly by requiring new equipment, there is also evidence pollution controls raised costs indirectly by increasing the complexity of the plant, which now required greater planning and longer construction times [11, 36]. We refer the reader to several studies that examine more thoroughly the contribution of pollution controls to plant costs versus other factors (most notably inflation, interest rates, and construction productivity) [32, 10, 11, 23, 36].

Other causes cited for the increase in costs are decreasing construction productivity and increasing construction times [11, 36], reversed economies of scale (plants got smaller after 1970), and diminished opportunities and incentives to improve plant construction due to design variation and principal-agent problems [12]. The relative contribution from each factor is unclear. For example, different sources report that pollution controls raised costs anywhere from 20-66% [11, 12, 37].

A common representation of technological improvement is the experience curve or performance curve, a plot of the cost versus cumulative production of some item. There is a large literature debating the value of experience curves as a model for technological improvement, which we review in Section 5.1. In Fig. 6 we show the experience curve for plant construction costs. We use two different measures of experience: the number of coal units built, and the total capacity installed. Note that the cluster of points on the right side of Fig. 5 is separated from the last point in 1987 by a gap due to missing data. However, in both plots of Fig. 6, this cluster of points joins up with those from the 1980s. This is consistent with experience curve hypothesis that costs are more correlated with changes in experience (as represented by cumulative production) than with time.

2.5.2. Capacity factor

The capacity factor (CF) is the amount of electricity produced divided by total potential production:

$$CF_t = \frac{\text{kWh production in year } t}{\text{kW capacity} \times 8760 \text{ hrs/year}}$$

The capacity factor measures the utilization of a plant's capacity, and is bounded between 0 and 1.

Electricity generation incurs large fixed costs from plant construction, making high capacity factor – high utilization of the plant's capital – desirable to spread costs over the greatest possible production. The capacity factor is largely determined in advance, by the choice to build a base load plant or peaking plant. Demand for electricity varies hourly and seasonally. Building a coal plant with a high enough capacity to meet peak demand would leave the plant underused during periods of low demand, effectively raising capital costs. The cheapest way to meet

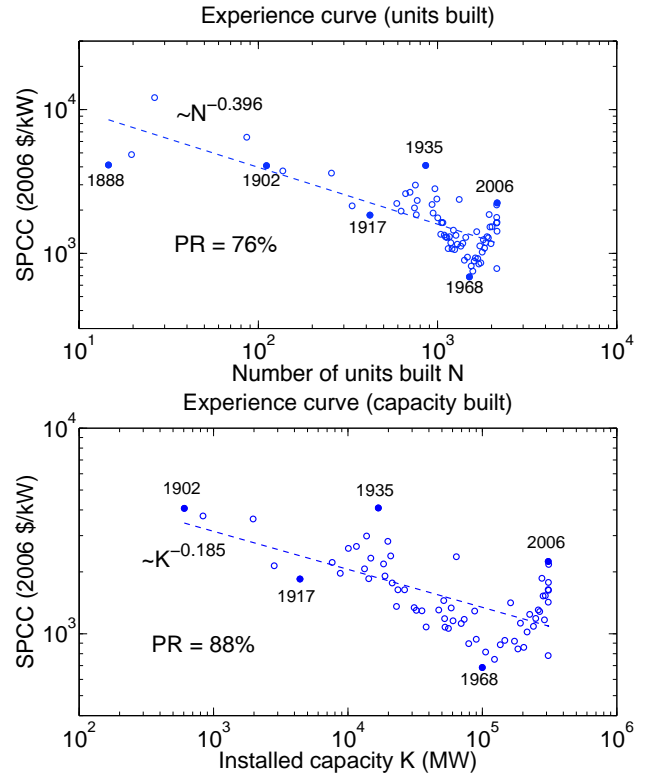


Figure 6: Experience curves for plant construction cost, using two possible experience measures. Experience curves are a popular framework for predicting future costs of technologies. Source: [33, 34, 35, 38, 39, 17]

varying electricity demand is with a combination of base load plants that have low operating cost (O&M plus fuel) and run at high output rate to cover the persistent portion of electricity demand, and peaking plants that are relatively inexpensive to build and cover the excess portion of demand unmet by the base load plant.

Modern coal plants usually serve as baseload plants and therefore tend to have high capacity factors, around .7 – .8 (Fig. 7), though over the last 50 years the capacity factor has varied between .50 and .82. Capacity factors before 1940 were much lower. This may be because early plants required frequent maintenance and few devices existed that required electricity, resulting in less consistent demand for electricity throughout the day. The first major use of electricity was for lighting, particularly street lighting, which was only necessary a few hours each day.

2.5.3. Interest rate

As mentioned earlier, we amortize the specific construction cost of plants (SC) to obtain the capital cost (CAP) using

$$CAP_t = \frac{SC_t \times CRF(r_t, n)}{CF_t \times 8760 \text{ hours/year}} \quad (3)$$

CRF is the capital recovery factor defined in Eq. (2). We take the plant lifetime n to be 30 years for amortization purposes. This is both conventional and matches




Figure 7: Capacity factor of coal plants. Dashed lines are linear interpolations. Coal plants are typically base load plants, serving the persistent part of electricity demand at high capacity factor. The growth of the capacity factor, indicating greater utilization of the plant’s capital, was a substantial factor reducing coal-fired generation costs. Source: [39, 20, 32, 38, 17]

the longest bond maturities, the period over which capital payments would be made.³ The resulting CRF for a given year was 1-2% higher than the annual interest rate for that year.

The effective interest rate used to calculate the capital recovery factor was the average return-on-investment (ROI) of the electric utility industry in each year (Fig. 8 inset.) The ROI is defined as the sum of annual interest and dividend payments made to investors divided by the industry’s gross plant value. This data was gathered from combined income statements and balance sheets for the electric utility industry as a whole [39, 41, 20]. The data is therefore not coal-specific, but we do not expect that interest rates charged to coal utilities would differ significantly from the average rate charged to utilities as a whole. The ROI varied between 4 and 8%.

Interest rates had an important but transient effect; they contributed to the rise in costs seen during the 1970s and 1980s.

2.5.4. Capital cost

The capital cost (CAP) given by Eq. 3 combines construction costs, plant usage, and interest rates to obtain the contribution of capital to total generation costs (Fig. 8 main axes.) Its shape resembles that of the specific construction cost of Fig. 5. The most important factors reducing capital costs were decreasing construction costs and increasing plant usage.

2.6. Total cost

Figure 9 shows the total generation cost of coal-fired electricity (TC), along with the three major cost compo-

³However, as a point of interest, it is quite possible that plant lifetimes are not 30 years, and moreover have changed with time. Longer plant lifetimes would effectively reduce electricity costs by spreading capital costs out over a longer period. One study found that most coal plants currently in use have lifetimes between 31 and 60 years, with an average of 49 [40].




Figure 8: The capital cost of coal plants. The capital cost accounts for the specific construction cost (Fig. 5), the capacity factor (Fig. 7), and the interest rate (inset.) This series resembles that of the specific construction costs but is steeper due to increases in capacity factor. The inset shows the historical return-on-investment of electric utilities, used as the effective interest rate for amortizing the construction cost. The Aaa corporate bond rate is shown for comparison. Source: [42, 39, 41, 20]




Figure 9: The total cost of coal-fired electricity. The top curve is the sum of the fuel, capital, and O&M curves below it.

ments. To check that the cost history constructed was approximately correct, we sought independent data to validate this series. To our knowledge, neither cost nor price data for coal-fired electricity exists, so we use the average price of electricity from all types of generation. Coal provided about half of all annual electricity production for the US throughout its history, so we expect the average price to be heavily influenced by the price of coal-fired electricity. In addition, we use historical data for transmission and distribution losses, taxes, and retained earnings (i.e. “post-cost” adjustments) to estimate the average *cost* of electricity to obtain a theoretically closer series for comparison. (These series can be seen in the Supporting Information.) The all-sources cost and price series described above, along with the reconstructed coal-fired cost, are shown in Fig. 10. These series compare favorably and validate the decomposition.

Between 1970 and 1985, costs increased due to a number of unrelated developments that happened nearly simultaneously. Coal prices increased from high oil prices,




Figure 10: The total cost of coal-fired electricity shown in Fig. 9, along with the average price and cost of electricity from all generating sources. The cost from all sources, which largely consists of coal-fired generation, was used as a point of comparison to validate our built-up cost for coal-fired electricity. The former is expected to be somewhat higher because it includes electricity from more expensive sources.

anticipation of strikes, and increasing wages; O&M costs increased because of added pollution controls; and plant construction costs increased from added pollution controls, diminishing productivity in the construction sector, and high inflation-driven interest rates.

In the next three sections we analyze the data presented above. In Section 3 we look at the long term cost trends of each variable. In Section 4 we look at the short term variation caused by each variable.

3. Analysis of cost trends

We analyze the trends in total generation costs in two ways. First, we show the contribution of each major cost component to total generation costs (TC). Second, we show the contribution of each variable to changes in the total generation cost.

The major contributors to total generation cost are the fuel and capital components. The contribution from fuel has usually been 50-65%, and that from capital 25-40%.




Figure 11: The share of total generation costs contributed by each component. Fuel and capital have traded off in dominance of total generation costs over time, and are currently close in size.

However, during two periods this ordering failed – from about 1925 to 1939, when the roles of fuel and capital swapped places, and from 1984 to the present, where they contribute about evenly. The contribution of O&M costs appears to have been relatively steady around 5-15% during the whole history (Fig. 11), though as noted earlier O&M costs before 1938 are not based on data but inferred from fuel costs (Fig. 4). This estimated breakdown is similar to ones given in other sources [32, 36].

The second way we break down total costs is to decompose the *change* in total cost contributed by each variable. That is, we calculate the change ΔTC_i in the total generation cost caused by each variable i . These contributions from individual variables sum up to the total change ΔTC :

$$\begin{aligned} \Delta TC &= \Delta TC_{\text{OM}} + \Delta TC_{\text{FUEL}} + \Delta TC_{\text{CAP}} \\ &= \Delta TC_{\text{OM}} + \Delta TC_{\text{COAL}} + \Delta TC_{\text{TRANS}} + \Delta TC_{\eta} \\ &\quad + \Delta TC_{\rho} + \Delta TC_{\text{SC}} + \Delta TC_r + \Delta TC_{\text{CF}} \\ &= \sum_i \Delta TC_i \end{aligned}$$

The method for calculating each ΔTC_i is a generalization of Nemet’s method [14]. Our generalization is based on partial derivatives and is described in Appendix A.

Table 2 displays the results of this second decomposition during two periods, 1902–1970 and 1970–2006. The change in the generation cost ΔTC is given in the first row of the table; e.g. the generation cost dropped 27 ¢/kWh between 1902 and 1970 and increased 1.8 ¢/kWh between 1970 and 2006. The next three rows indicate the contributions to this change coming from the three major cost components: ΔTC_{OM} , ΔTC_{FUEL} , and ΔTC_{CAP} . These contributions sum up to ΔTC . The next eight rows decompose these contributions still further. Together these eight contributions also sum up to the total change in generation cost. Appropriate combinations of them will also sum to equal the contributions from OM, FUEL, and CAP.

The last two columns of the table give the results in percentage terms. The percent change in generation cost effected by variable i between years t_1 and t_2 is

$$\% \text{ of change caused by } i = 100 \times \frac{\Delta TC_i(t_1, t_2)}{\Delta TC(t_1, t_2)} \quad (4)$$

where $\Delta TC_i(t_1, t_2)$ denotes the change caused by variable i between years t_1 and t_2 . By construction all percent contributions sum to 100%. An implication of Eq. (4) is that variables which *oppose* a given change in generation cost appropriately have negative percent contributions. For example, while in the first period total generation costs *dropped* about 27 ¢/kWh (representing 100% of the total change), coal prices nevertheless increased slightly, and by themselves would have *increased* generation costs by .65 ¢/kWh (thus representing -2.4% of the total change.)

During the first period, two factors stand out most: increasing capacity factor, responsible for 26% of the decrease in costs, and improving efficiency, responsible for

Table 2: Decomposition of the change in cost of coal-fired electricity. The first two columns indicate the dollar amount of cost changes contributed by the individual variables of equation (1). The last two columns indicate what percentage of the change each variable is responsible for. Negative percent contributions indicate that a variable opposed the change that was actually realized.

Variable i	Effect on generation cost, ΔTC_i		% of ΔTC caused by i	
	1902 – 1970	1970 – 2006	1902 – 1970	1970 – 2006
TC	-27.24 [†] ¢/kWh	1.764* ¢/kWh	100.0 %	100.0 %
OM	-2.94	0.171	10.8 %	9.7 %
FUEL	-14.08	-0.104	51.7 %	-5.9 %
CAP	-10.22	1.697	37.5 %	96.2 %
OM	-2.94	0.171	10.8 %	9.7 %
COAL	0.65	-0.088	-2.4 %	-5.0 %
TRANS	-0.71	-0.286	2.6 %	-16.2 %
η	-15.04	0.021	55.2 %	1.2 %
ρ	1.01	0.249	-3.7 %	14.1 %
SC	-3.38	1.544	12.4 %	87.5 %
r	0.20	0.235	-0.8 %	13.3 %
CF	-7.03	-0.081	25.8 %	-4.6 %

[†] This is a 90% drop from the 1902 generation cost.

* This is a 57% rise from the 1970 generation cost.

55%. Following distantly are the specific capital cost (12%) and O&M cost (11%). The fuel component as a whole was responsible for 52% of the decrease in cost, and the capital component for 38%.

Surprisingly, the construction cost was only responsible for 12% of the change, despite decreasing by a factor of 5.4 during this time. However, it is important to realize that the magnitude of individual changes cannot be considered independently of other changes. If other cost-decreasing changes had not occurred simultaneously, the contribution of specific capital cost would have been much greater. Equally, if construction costs had not come down, it would have held up progress of the technology and changes in other factors would have been less important. Note that this is not a bug of the decomposition method used, but a fact occurring for any decomposition because changes in variables cannot be considered independently of other variables.⁴

During the second period, increases in plant construction costs contributed dramatically to the *increase* in generation cost. The capital component as a whole rose sufficiently by itself to double total costs. The O&M cost made a much smaller contribution to the increase, while a small change in fuel costs mitigated the total change in generation costs somewhat.

⁴Consider for example, some quantity y which depends on three variables x_1 , x_2 , and x_3 :

$$y(x_1, x_2, x_3) = x_1 x_2 + x_3$$

Suppose that x_2 changes by 100%, and x_3 changes by 5%. Does this mean that x_2 's contribution to the change is greater? Not necessarily; if x_1 is sufficiently small, then the change in the *product* $x_1 x_2$ may be tiny compared with the 5% change in x_3 . Thus, there are important "cross effects" of variables, which cannot be avoided and are a general fact of decomposing changes.

So far, we have studied the long term evolution of each variable and its effect on the generation cost. In addition to long term trends, though, some variables show significant short term variation. In the next section, we examine the influence of this variation on the generation cost.

4. Analysis of cost variation

We now look at the influence of short time (10 year) variations on the generation cost. We can do this using the same decomposition technique used above. First, we calculate ΔTC_i for each variable i between two given years t_1 and t_2 , as before. Then, instead of comparing ΔTC_i to ΔTC , we compare it to the value of TC in year t_1 :

$$\% \text{ variation caused by } i = 100 \times \frac{\Delta TC_i(t_1, t_2)}{TC(t_1)}.$$

This measures how much the change in i actually raised the generation cost between years t_1 and t_2 . Note that the previous section asked "how much of the generation cost change ΔTC does ΔTC_i represent?," regardless of the actual size of ΔTC . Now we are asking whether ΔTC_i actually changed the cost significantly.

Rather than showing the variation in generation cost (TC) contributed by all the variables presented in the previous section, we focus on the three largest contributors of variation: the price of coal (COAL), the specific construction cost of plants (SC), and the capacity factor (CF). Fig. 12 shows the influence these variables had on the generation cost. The height of the bar at each year t shows the value of $\Delta TC_i(t - 10, t)/TC(t - 10)$; i.e. the size of the price change caused by variable i over the previous 10 years. Thus, Fig. 12 uses a sliding window which considers changes over every possible 10-year period.

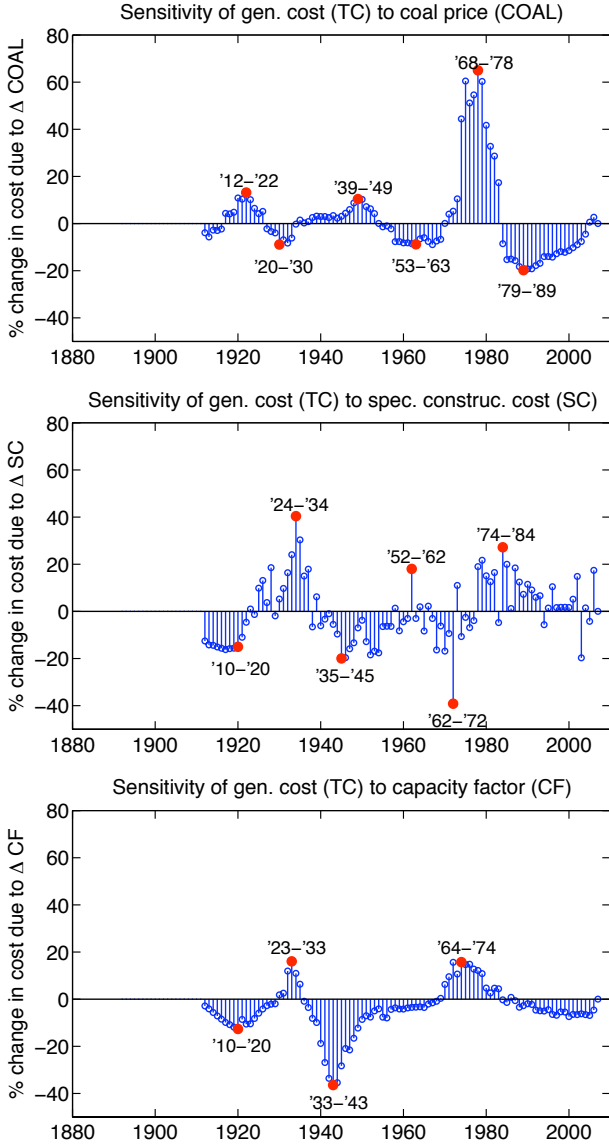


Figure 12: The variation in total generation cost, TC, caused by changes in the coal price, COAL (top), specific construction cost, SC (middle), and capacity factor, CF (bottom). Heights of vertical lines estimate how large a change in the total generation cost was caused by the change in the coal price over the preceding 10 years. Time intervals with historically large changes are highlighted in red. These variables caused the largest cost changes out of those given in equation (1).

The largest increase over any 10-year span came from coal price changes between 1968 and 1978, when coal prices alone caused generation costs to increase by 64.3%. The total increase in generation cost during this same period was 131.1%. The remainder of the increase was driven mostly by increasing O&M and capital costs.

Nevertheless, the construction cost and the capacity factor also contributed significant variation. Whereas coal price variation has tended to be smooth, with changes from consecutive years reinforcing each other, variation from construction costs is more noisy from year to year




Figure 13: Number of coal units produced (top), and geometric average capacity of new units (bottom). Source: [38]

(although still following a long term trend.) Unlike coal, which is more or less the same product every year, each year's batch of new coal plants has unique characteristics that may make it more or less expensive than those of neighboring years. Thus, some of the spikes seen in the second plot of Fig. 12 are due to two batches of plants 10 years apart that happen to have significantly different average costs.

Finally, although the capacity factor also contributed large variation, it is mostly an exogenous factor. It depends on industry-wide trends, rather than on any factors unique to coal-fired electricity, as mentioned before.

5. Extrapolation

We build a simple model for extrapolating generation costs. The goal of our extrapolation is to better understand the constraints on the baseline generation costs of coal-fired electricity. By *baseline costs* we mean the cost of generation without carbon capture and storage (CCS). Much discussion of future costs of coal-fired electricity focuses on the costs of plants equipped for CCS and the expected improvement over time [43]. Our focus instead will be to understand how the cost of the underlying technology may evolve, without modifications required by CCS, since historical data on CCS is not available. Furthermore, as noted earlier, in order to understand the impact of any market-based effort to reduce carbon emissions, it is nec-

essary to understand how the generation costs for existing, well-established technologies will evolve in comparison to alternative technologies.

In the following model, we treat the key variables driving fuel and capital costs in different ways. The driving variable behind the capital cost has been the specific construction cost, while the driving variable behind the fuel cost since efficiencies stopped improving has been the price of coal. The specific construction cost has tended to follow steady long term trends. In contrast, the price of coal seems to vary randomly, with no clear trends. As discussed in the next section, we hypothesize that there is a fundamental distinction between the behavior of a commodity and a technology, which is relevant to understanding the factors that influence coal-fired electricity.

5.1. Commodities vs. pure technologies

In this section we present the conjecture that the prices of commodities and technologies evolve in fundamentally different ways. The meaning of these terms will be clarified momentarily, but the relevance to coal-generated electricity is that coal is essentially a commodity, whereas the construction cost of a plant is closer to (though not purely) a technology. This discussion motivates the simple forecasting model we develop in the remainder of the section.

A commodity is a raw material or primary agricultural product that can be bought and sold. A standard assumption in the theory of finance is that markets are efficient; roughly speaking, this means it should not be possible to make consistent profits by arbitrage of commodities using simple strategies. If the price of coal were too predictable using simple methods, such methods would become common knowledge, and the buying and selling activity of speculators would affect prices in a way that would destroy their predictability. More specifically, one expects prices to follow a random walk to some approximation. Furthermore, the activity of speculators would bound the long-term rate of growth or decline of the price. Otherwise it would be possible to make unreasonable profits – more reliably than could be made in alternative investments – by either buying coal and hoarding it or by short selling coal.⁵

A pure technology is a body of knowledge, such as knowledge of a manufacturing process. We will say a product behaves like a pure technology when accumulated knowledge, reflected in changes to the technology’s design, is a greater determinant of its cost evolution than speculation. For example, the fuel component of coal-fired plants were initially more pure technology-like. Changes to boiler and plant design [44] resulted in growth of thermal efficiency and pure technology-like cost evolution in the fuel component (which depends on the thermal efficiency.) Although the input cost of coal could and did

⁵Other ways to speculate on the price of coal are to buy or sell coal companies and to buy or sell land containing coal deposits.

change, its changes had less impact on the fuel component cost than these pure design changes. In the present day, with many of the efficiency improving design changes already exploited, changes to input costs determine changes in the fuel component to a greater extent, making the fuel component more commodity-like.⁶

Commodities and pure technologies are ends of a spectrum, and probably no real products are completely one or the other. The specific construction cost, for example, depends on construction technologies – that knowledge of construction process and design which separates efficient uses of materials and labor input from inefficient uses. However, specific construction cost also depends on material commodities, such as steel. Thus specific construction cost is partly a technology and partly a commodity. In section 5.5 we choose to model it as a pure technology because we think this is its dominant character. Likewise coal is not completely a commodity, since knowledge of coal mining methods also affect its price, in addition to the speculation mechanisms mentioned above. A third example of this mixed composition is a photovoltaic system. The price of a photovoltaic (PV) system depends in part on commodities, such as metals and silicon. Increases in the price of silicon, for example, caused increases in the cost of PV cells in 2005-2006. Nonetheless, silicon is far from a raw material – processing dominates the cost of mono and polycrystalline silicon, rather than commodity input costs. Thus, while the price of PV cells is partially commodity-driven, it may behave more like a pure technology than a commodity. Its historical behavior suggests this is true (see the silicon price time series in Nemet [14]), though this possibility deserves further investigation.

5.2. Modeling a commodity

The price of a commodity is usually modeled using time series methods. Time series models form a class of nested models that can be arbitrarily complex, but the simplest model and conventional starting point is an autoregressive model of order 1 (AR(1).)

The AR(1) model is defined by the equation

$$p_t = \gamma p_{t-1} + \mu + \varepsilon_t.$$

Here p_t is the logarithm of the price in year t , ε_t is a random variable, and γ and μ are parameters. The noise

⁶It’s also reasonable to wonder the extent to which coal itself is a pure technology versus a commodity. This question is further complicated by the fact that coal is influenced by scarcity factors as well as technological factors. Anecdotal evidence makes it clear that both of these are influential. Both of these factors can, and probably did, cause changes in the cost/price of coal. (Although it is a bit interesting that they should stay so balanced with each other over a 130 year period. We say this noting that the dramatic rise in the 70s was largely due to exceptional circumstances having little to do with scarcity or technology factors.) Nevertheless, the theory of finance bounds the *anticipated* long-term changes that can occur. This allows for unanticipated changes that may have any size. It also allows for anticipated long term changes that are sufficiently slow that they are not worth taking advantage of.

term ε_t is assumed to be uncorrelated in time and normally distributed with variance σ_ε^2 . The three parameters γ , μ , and σ_ε , combined with an initial condition $p(0)$, determine a given AR(1) process.⁷

The parameter γ determines the long-term behavior of the process. When $\gamma = 1$, the process is a random walk with drift, and when $\mu > 1$ the mean of the distribution of future prices grows without bound. When $\gamma < 1$, the process is stationary and the distribution of future prices asymptotically approaches a fixed width. The closer γ is to 1, the wider the limiting distribution becomes. For $\gamma = 1$, the parameter μ generates drift in the random walk, but for $\gamma < 1$ it sets a non-zero mean value around which the price moves.

5.3. Modeling a technology

Extrapolating a technology requires using some effective model for its behavior. We will discuss two common effective models for extrapolating technologies.

The cost of a technology is often modeled in terms of an experience curve, a plot of its cost versus its cumulative production. The motivation for plotting cost against cumulative production often uses the following chain of reasoning: cumulative production measures “experience” with a technology; as a firm or industry gains experience it makes improvements to the technology, which cause cost reductions; therefore there may be a simple, predictable relationship between cumulative production and cost. Alternatively, one can argue that cumulative production is an indicator of profit-making potential, which drives the level of effort directed at improving a given technology [45].

In fact, a simple relationship is frequently observed: empirical experience curves often appear to obey power laws. There is a large literature attesting to this regularity [46, 47]. Partly for this reason, the experience curve, combined with the assumption of a power law form, has become the prevailing method for extrapolating future costs [48]. The power law functional form can also be derived from theoretical models [49, 50].

We find it useful to distinguish the *data* from any particular *hypothesis* about its shape. We therefore use the term experience curve to mean the plot of cost versus cumulative production, whatever its shape. We reserve the term Wright’s law for the hypothesis that an experience curve should follow a power law. While we acknowledge that many researchers implicitly include the power law functional form with the definition of an experience curve, we prefer to refer to the data in a theory-neutral way.

Alternatively, the cost of technologies has been modeled as decreasing exponentially with time; this evolu-

tion can be viewed as a generalization of Moore’s law.⁸ Moore’s law (an *exponential* decrease in cost with *time*) and Wright’s law (a *power law* decrease in cost with *cumulative production*) are not necessarily incompatible. They could simultaneously hold when cumulative production grows exponentially with time.⁹

Two kinds of criticisms have been brought against the program of plotting cost data and fitting it to curves.¹⁰ One is that it buries too much detail about the processes driving cost reductions [45]. Another criticism questions the predictive power of different functional forms, since they have not been rigorously compared against each other. Further progress in this area requires both careful comparison of the predictive ability of different effective models and study of the causes underlying observed relationships. Nagy et al. have recently compared the historical prediction accuracy of various functional forms proposed by researchers to describe cost evolution [51], including Wright’s law and Moore’s law.

Finally, we note that any mathematical description of a technology’s cost evolution is likely to fail when the technology itself substantially changes. A famous example is the Model T Ford, which dropped smoothly in cost from its introduction in 1909, when Henry Ford announced he would make a car that the common could afford, to 1929, when he ceased production. During this period the cost of Model T’s follow Wright’s law quite well [53]. After that Ford produced other models with better performance and, not surprisingly, higher costs. This is an important point; in general an extrapolation may only hold when the product faces fixed performance criteria. When a product is called upon to meet criteria it did not previously meet, such as lower pollution or higher safety, its cost evolution may show discontinuous behavior.

5.4. Capital and O&M costs

The principal driving variable for capital cost is specific capital cost. As already discussed above, and based on the empirical evidence presented in Section 2.5.1, it is reasonable to model specific cost as a technology. We have modeled the specific cost both with Wright’s law and with

⁷“Order one” refers to the number of earlier prices the model regresses the current price onto. An AR(3) model, for example, would be specified as $p_t = \gamma_1 p_{t-1} + \gamma_2 p_{t-2} + \gamma_3 p_{t-3} + \mu + \varepsilon_t$.

⁸See Nagy et al. [51]. Moore’s law originally states that the number of transistors per integrated circuit doubles about every 2 years on average. Basically, this regularity about transistor density can be restated in terms of the cost of processing power; this transformation shows that Moore’s law implies the cost of processing power decreases exponentially over time. Thus it is converted into a statement about cost rather than transistor density. The new statement of Moore’s law is equivalent to the original one, but has the advantage of being expressed in general terms comparable across technologies (i.e. cost versus time.) Moreover, it shows that transistors join up with other technologies whose costs appear to drop exponentially with time. Thus we use Moore’s law simply to mean the hypothesis that cost decreases exponentially with time.

⁹See, for example, Sahal [52].

¹⁰These criticisms have often been directed against experience curves in particular, but could apply to any low-dimensional model of cost evolution.

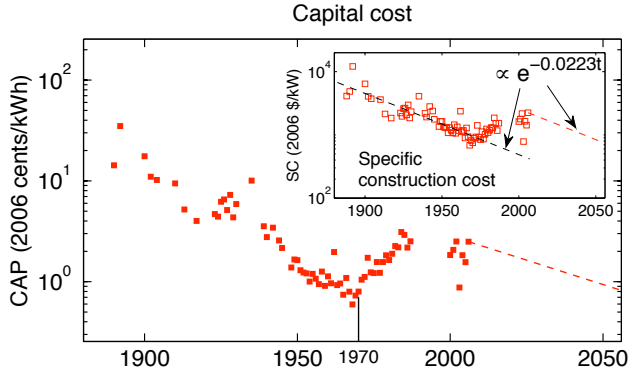


Figure 14: Extrapolation of capital costs 50 years into future. Inset shows a low-cost scenario for specific construction cost based on assumptions given in the text. Main axes show the resulting capital cost after accounting for capacity factor and interest rate. These variables are assumed to stay fixed at their 2006 values.

the generalized version of Moore's law. We find the end results yield similar predictions, and so choose to use an exponential in time because it is simpler.

One of the problems we confront in modeling the specific capital cost is the break in the decreasing trend that occurs roughly in 1970, as shown in Figure 14. As already discussed in Section 2.5.1, roughly speaking the specific construction cost generally dropped from 1902 - 1970, increased until about 1990, and has been roughly flat or slightly decreasing since then. Pollution controls have been proposed as the dominant cause of the increase in prices that occurs after 1970, though this is somewhat controversial. To make an extrapolation model we are forced to make a choice as to how to treat the period of rising costs, and what its effect will be going forward.

Our main point in this paper is to show that, even under an optimistic point of view in which the construction cost of coal plants substantially decreases, the price of coal will dominate the cost of coal-fired electricity, and will impose a fluctuating floor below which it is unlikely to penetrate. Thus, in order to be as optimistic as is reasonably possible about the future behavior of coal-fired electricity prices, we postulate a *low cost scenario* in which we assume that after 2006 the specific construction cost reverts to the same rate of exponential decrease that it had prior to 1970.

To implement this scenario we fit an exponential curve to the historical data only from 1902 - 1970, as shown in Figures 8 and 14. Then we generate an exponential extrapolation having this same slope, but multiplied by an appropriate factor to have it start from the 2006 level of specific construction cost (Fig. 14 inset.) We follow a similar procedure for O&M costs. All other variables used in our cost equation are assumed to maintain their 2006 values. The resulting extrapolation provides a crude but simplified way of estimating potential baseline generation costs on the basis of historical experience.

The inset of Fig. 14 shows the specific construction

cost under the assumption that plants return to their pre-1970 rate of improvement, while the main axes show the resulting evolution of the amortized capital cost after accounting for the 2006 values of interest rate and capacity factor. Under the low cost scenario, the capital cost would be reduced 67% by 2056 to 0.82 ¢/kWh. This would still place it higher than the historical low of 0.56 ¢/kWh that occurred in 1968. (We omit the plot of the O&M extrapolation because it is less important, but we include its contribution in the total generation cost shown in Fig. 16.)

5.5. Fuel costs

Coal prices appear to behave quite differently than specific construction cost. During the course of its history the price of coal fluctuates in a random-looking manner, with no clear trend toward either increasing or decreasing values. We hypothesize that this is to be expected based on the fact that coal is a commodity.

As explained in Section 5.2, the AR(1) model is the conventional starting point for analyses of price time-series data. We fit the historical coal price to the AR(1) model using least squares regression. The parameter values obtained are $\gamma \approx .956 \pm .053$, $\mu \approx .144 \pm .174$, and $\sigma_\epsilon \approx .080$. Note that γ has an error interval that includes $\gamma = 1$. Since $\gamma = 1$ has qualitatively different behavior than $\gamma < 1$, it is important to know the likelihood that the underlying process actually has $\gamma = 1$ and is therefore a random walk. We apply various Dickey-Fuller tests of the null-hypothesis that the underlying process is a random walk, and find that in no test can this possibility be rejected. Thus, our hypothesis is validated in two different ways: First, to within statistical error $\gamma \approx 1$, suggesting a random walk, and second, to within statistical error the drift rate $\mu \approx 0$, suggesting no long term trend in coal prices.

The inset of Fig. 15 shows 10 simulations of the coal price 50 years into the future, evolving according to an

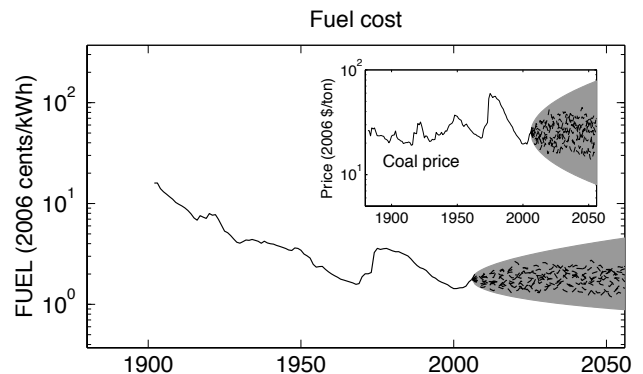


Figure 15: Simulation of fuel cost 50 years into the future. The inset shows the price of coal, the main axes the fuel cost after accounting for transportation costs, coal energy density, and plant efficiency. Dispersion in the simulated runs indicates the range of uncertainty for an AR(1) process. The boundaries of the filled-in region show a two-standard deviation envelope for the distribution of a random walk, indicating the range of uncertainty for a random walk.

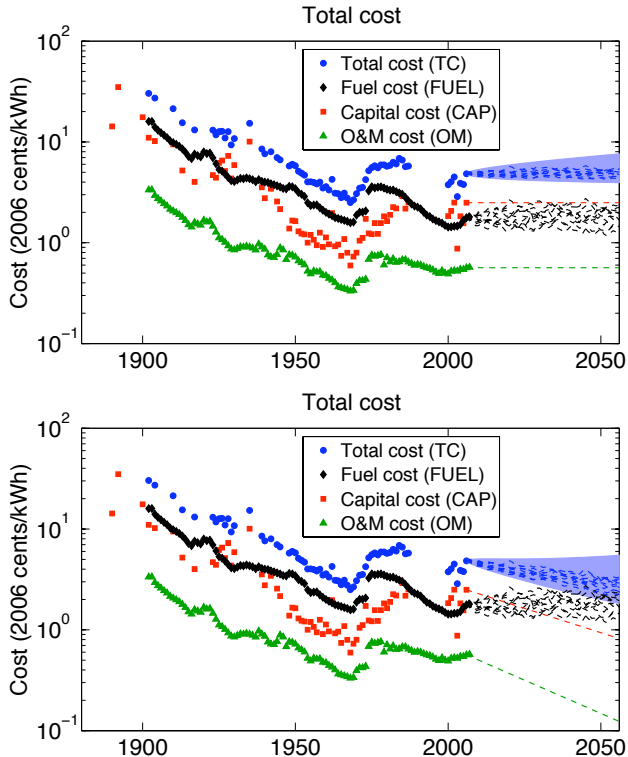


Figure 16: The lower figure shows the estimated low-cost scenario of coal-fired generation costs, and the resulting sensitivity to the coal price. Dispersion in the simulated runs indicates the range of uncertainty if coal prices undergo an AR(1) process. The boundary of the filled-in region indicates the range of uncertainty if coal prices undergo a random walk. In a low-cost scenario where plant costs decrease at the pre-1970 rate, generation costs would be expected to drop 42% by 2056, though preeminence of fuel costs makes the size of the drop uncertain and changing with time due to fluctuating coal prices. For comparison, the top figure shows the smaller uncertainty and fluctuation that would result if plant construction and O&M costs remain at their current values. The relatively flat evolution of fuel costs follows from the assumption that efficiency realizes no dramatic increases.

AR(1) process with the fitted values for γ , μ , and σ_ε . The main axes show the resulting evolution of the fuel component cost after accounting for the 2006 values of transportation costs, coal energy density, and plant efficiency. The shaded region indicates the envelope of uncertainty in the evolution of the price. Because a random walk could not be rejected, the envelope is that produced by a random walk. While the generated series do not resemble the historical series in all respects—they have more high-frequency variation, for example—they show an amount of dispersion similar to that of the original series.

5.6. Extrapolation summary

To summarize our extrapolation model that we developed above, we make 4 assumptions:

1. Fuel costs evolve according to an AR(1) process.
2. Plant construction costs follow their average 1902-1970 slope.

3. O&M costs follow their average 1902-1970 slope.
4. All other quantities remain fixed at their 2006 values.

These assumptions lead to the predictions given in the next section.

5.7. Extrapolation results

Based on this extrapolation, generation costs could drop up to 42% between 2006 and 2056 under a low-cost scenario for plant construction costs (Fig. 16.) By the end of this period, plant costs would contribute 29% of the total generation costs. Dropping plant costs would cause fuel to become the dominant cost, contributing two-thirds of the generation costs. The dominance of fuel costs would in turn lead to greater sensitivity of generation costs to fluctuations in the price of coal. The 10 simulations in Fig. 16 show the approximate range of extrapolations when coal prices follow an AR(1) process. The dotted lines show the approximate range of extrapolations when coal prices follow a random walk. Note that this range only accounts for uncertainty in the coal price. Uncertainty in other variables would increase the range.

This scenario assumes no dramatic change in thermal efficiency. As mentioned in Section 2.3.4, efficiency has remained flat the last 50 years for several reasons. Moreover, we find it would take a substantial increase in efficiency to yield a substantial change in these results. For example, if the average efficiency increased from the 33% to 43% at the same time as the assumed decrease in plant costs occurred, the size of the potential drop in cost goes from 42% to 51%. Fuel would then represent 60% of generation costs. Note however that gains in efficiency usually come at the expense of increased plant costs, and an increase in the average efficiency would require a substantial upgrade of all plants, not just high performers.

6. Conclusions

We make several methodological advancements in this paper. We consider total generation costs (decomposed into components) rather than costs of single components alone, use data over a relatively long time span (over a century and covering the lifetime of the industry), and use a physically accurate bottom-up model of costs. As part of this decomposition, we model coal prices as a random walk while modeling construction costs and O&M costs as an improving technology.

We find evidence that the fuel and capital costs of coal-fired electricity evolve differently due to different behaviors in coal prices and plant construction costs. Coal prices have fluctuated with no trend up or down, while plant construction costs have followed long-term trends. The behavior of coal prices is consistent with the facts that coal is a freely traded commodity. Plant construction cost may follow a pattern of long term reduction because it is only weakly influenced by these effects and is able to realize improvements typically seen in technologies over time.

Such a difference in the behavior of coal prices and plant construction costs would yield different behavior for the fuel and capital cost components. Although historically both the fuel and capital components improved at similar rates, the main driver of improvement in fuel costs – thermal efficiency – has been unchanged since the early 1960s. Without a substantial improvement in the average thermal efficiency, the main driver of change in fuel costs would be the coal price. Under a low-cost scenario where capital costs were driven down by improving plant construction costs, the total generation cost would decrease about 40% over the next 50 years. Coal prices, in contrast, are statistically neither decreasing nor increasing, and so provide a statistically fluctuating floor on the overall generation cost, with no clear long-term trend.

Although all the analysis in this paper is specific to coal, we hypothesize that a similar analysis might apply to other fuel-based sources of electricity, such as natural gas. In fact, because natural gas and oil are both traded on exchanges, with standardized futures contracts, it is easier to speculate in them than it is in coal. Thus we would expect that the commodity part of our model will apply to them even more strongly. This would suggest that the major non-renewable sources of electricity should all hit a noisy floor, similar to the one we have postulated here for coal.

7. Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 0624351 (this is the Financial Markets grant no). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. We thank Margaret Alexander and Tim Taylor for substantial help obtaining source materials. We thank Arnulf Grübler for helpful conversations and suggestions. We thank Charlie Wilson for several valuable conversations and extensive feedback on an early draft, and Dan Schrag for a challenge that led to the idea for this paper.

A. Change decomposition

Consider the following problem. We have a function $f = f(x, y)$, and during some period of time f changes as a result of simultaneous changes to x and y . We want to know how much of the change to f each variable is “responsible for.”

To be more precise, let Δf be the change in f . We would like to decompose Δf into 2 terms, corresponding to the change contributed by each variable:

$$\Delta f = \Delta f_x + \Delta f_y,$$

where Δf_i denotes the change in f resulting from the change in i . A way to do this decomposition is suggested

by the calculus identity

$$df = \frac{\partial f}{\partial x} dx + \frac{\partial f}{\partial y} dy.$$

The trick is to generalize this expression to finite rather than infinitesimal changes. More to the point, we need to be able to take common combinations of variables – e.g. products, quotients – and express changes to these combinations in ways appropriate for finite differences.

For example, consider a product of two variables, $f(x, y) = xy$. The differential f is

$$df = x dy + y dx$$

suggesting the decomposition for finite differences is $\Delta f = x\Delta y + y\Delta x$. However, the correct expression is

$$\Delta f = x\Delta y + y\Delta x + \Delta x\Delta y.$$

In addition to the two expected terms, a third *cross term* containing both differences appears. As the differences become small, the cross term will vanish more quickly than the other terms, recovering the calculus limit $df = x dy + y dx$. However, for finite differences, the cross term potentially introduces a large residual if ignored.

In order to decompose the change in f into just 2 pieces, we evenly split the cross term into the other terms:

$$\begin{aligned} \Delta f &= x\Delta y + \frac{1}{2}\Delta x\Delta y + y\Delta x + \frac{1}{2}\Delta x\Delta y \\ &= (x + \frac{1}{2}\Delta x)\Delta y + (y + \frac{1}{2}\Delta y)\Delta x \end{aligned}$$

The first term may be interpreted as the change in f due to change in y ; the second as the change in f due to change in x . We therefore have our desired decomposition for the case of products of 2 variables:

$$\begin{aligned} \Delta f_x &\equiv (y + \frac{1}{2}\Delta y)\Delta x \\ \Delta f_y &\equiv (x + \frac{1}{2}\Delta x)\Delta y \end{aligned}$$

which by construction has the desired property $\Delta f = \Delta f_x + \Delta f_y$. Similar rules can be derived for products of 3 or more variables, for quotients, and other expressions.

References

- [1] World bank development indicators database (2008).
- [2] Key world energy statistics, Tech. rep., International Energy Agency (2008).
- [3] R. Kumiya, Technological progress and the production function in the united states steam power industry, *The Review of Economics and Statistics* 44 (2) (1962) 156–166.
- [4] M. Nerlove, Returns to scale in electricity supply, in: *Measurement in economics: studies in mathematical economics and econometrics*, Stanford University, Stanford, 1963.
- [5] T. G. Cowing, Technical change and scale economies in an engineering production function: The case of steam electric power, *The Journal of Industrial Economics* 23 (2) (1974) 135–152.

- [6] L. R. Christensen, W. H. Greene, Economies of scale in U.S. electric power generation, *The Journal of Political Economy* 84 (4) (1976) 655–676.
- [7] D. A. Huettner, J. H. Landon, Electric utilities: Scale economies and diseconomies, *Southern Economic Journal* 44 (4) (1978) 883–912.
- [8] J. F. Stewart, Plant size, plant factor, and the shape of the average cost function in electric power, *The Bell Journal of Economics* 10 (2) (1979) 549–565.
- [9] R. A. Nelson, M. E. Wohar, Regulation, scale economies, and productivity in steam-electric generation, *International Economic Review* 24 (1) (1983) 57–59.
- [10] F. M. Gollop, M. J. Roberts, Environmental regulations and productivity growth: The case of fossil-fueled electric power generation, *Journal of Political Economy* 91 (4) (1983) 654–674.
- [11] P. L. Joskow, N. L. Rose, The effects of technological change, experience, and environmental regulation on the construction cost of coal-burning generating units, *Rand Journal of Economics* 16 (1) (1985) 1–27.
- [12] M. J. McCabe, Principals, agents, and the learning curve: The case of steam-electric power plant design and construction, *The Journal of Industrial Economics* 44 (4) (1996) 357–375.
- [13] J. J. Hisnanick, K. O. Kymn, Modeling economies of scale: the case of US electric power companies, *Energy Economics* 21 (1999) 225–237.
- [14] G. F. Nemet, Beyond the learning curve: factors influencing cost reductions in photovoltaics, *Energy Policy* 34 (2006) 3218–3232.
- [15] A study of coal prices, Tech. Rep. NP-20948, Council on Wage and Price Stability, Washington, DC (1976).
- [16] R. S. Manthy, J. R. Tron, *Natural Resource Commodities—A Century of Statistics*, Johns Hopkins, Baltimore, 1978.
- [17] Annual energy review, Tech. Rep. DOE/EIA-0384(2006), Energy Information Administration (2006).
- [18] S. H. Schurr, B. C. Netschert, *Energy in the American economy, 1850-1975: an economic study of its history and prospects*, Johns Hopkins, Baltimore, 1960.
- [19] Mineral yearbook, Tech. rep., US Geological Survey, Washington, DC (various issues 1932-1976).
- [20] *Historical Statistics of the Electric Utility Industry through 1992*, Edison Electric Institute, Washington, DC, 1995.
- [21] Encyclopaedia britannica (2009) [cited September 19, 2009]. URL <http://www.britannica.com/EBchecked/topic/615316/unit-train>
- [22] Steam-electric plant construction cost and annual production expenses for 1947, Tech. rep., Federal Power Commission (1947).
- [23] R. F. Hirsh, *Technology and Transformation in the American Electric Utility Industry*, Cambridge, 2002.
- [24] C. E. Neil, Entering the seventh decade of electric power, *Edison Electric Insitute Bulletin* 10 (9).
- [25] Steam-electric plant construction cost and annual production expenses, Tech. Rep. FPC S-154, Federal Power Commission (1961).
- [26] Steam-electric plant construction cost and annual production expenses, Tech. Rep. DOE/EIA-0033(78), Energy Information Administration (1978).
- [27] Thermal-electric plant construction cost and annual production expenses, Tech. Rep. DOE/EIA-0323(79), Energy Information Administration (1979).
- [28] A. Beamon, T. J. Leckey, Trends in power plant operating costs, Tech. Rep. EIA/DOE-0607, Energy Information Administration (1999). URL <http://www.eia.doe.gov/oiaf/issues/opctbl2.html>
- [29] U.S. electricity production costs and components (1995-2007), Tech. rep., Nuclear Energy Institute. URL <http://www.nei.org/resourcesandstats/documentlibrary/reliableandaffordableenergy/graphicsandcharts/uselectricityproductioncostsandcomponents/>
- [30] S. Yeh, E. S. Rubin, A centurial history of technological change and learning curves for pulverized coal-fired utility boilers, *Energy* 32 (10) (2007) 1996–2005.
- [31] P. Sporn, *Vistas in Electric Power*, Pergamon, Oxford, 1968.
- [32] C. Komonoff, *Power Plant Cost Escalation: Nuclear and Coal Capital Costs, Regulation, and Economics*, Van Nostrand, New York, 1981.
- [33] Historical plant cost and annual production expenses for selected plants, Tech. Rep. DOE/EIA-0455(82), Energy Information Administration (1982).
- [34] Historical plant cost and annual production expenses for selected plants, Tech. Rep. DOE/EIA-0455(87), Energy Information Administration (1987).
- [35] E. Shuster, Tracking new coal-fired power plants, Tech. rep., National Energy Technology Laboratory (2007).
- [36] B. L. Cohen, *The Nuclear Energy Option*, Plenum, New York, 1992.
- [37] G. M. Masters, *Renewable and Efficient Electric Power Systems*, John Wiley, Hoboken, 2004.
- [38] Platts UDI world electric power plants data base.
- [39] Census of electrical industries, Tech. rep., Census Bureau, Washington, DC (various issues 1902-1937).
- [40] H. Wynne, Tech. rep., Sanford C. Bernstein & Co. LLC (2008).
- [41] Statistics of privately owned electric utilities in the united states, Tech. Rep. FPC S-247, Federal Power Commission (1973).
- [42] Federal reserve statistical release H.15.
- [43] The future of coal, Tech. rep., Massachusetts Institute of Technology (2007).
- [44] B. . Veatch, *Power Plant Engineering*, Springer Science, New York, 1996.
- [45] W. C. Gavin Sinclair, Steven Klepper, What’s experience got to do with it? sources of cost reduction in a large specialty chemical producer, *Management Science* 46 (1) (2000) 28–45.
- [46] J. M. Dutton, A. Thomas, Treating progress functions as a managerial opportunity, *Acadamy of Management Review* 9 (2) (1984) 235–247.
- [47] P. Thompson, *Handbook of Economics of Technical Change*, Elsevier/North Holland, 2009, Ch. Learning By Doing.
- [48] J. D. Farmer, J. E. Trancik, Dynamics of technological development in the energy sector, in: J. P. Onstwedder, M. Mainelli (Eds.), *London Accord Final Publication*, 2007.
- [49] J. F. Muth, Search theory and the manufacturing progress function, *Management Science* 32 (8) (1986) 948–962.
- [50] J. McNerney, J. D. Farmer, S. Redner, J. E. Trancik, The role of design complexity in technology improvement.
- [51] B. Nagy, J. D. Farmer, J. E. Trancik, More than moore: Comparing forecasts of technological progress (in preparation).
- [52] D. Sahal, *Patterns of Technological Innovation*, Addison-Wesley, London, 1981.
- [53] W. J. Abernathy, K. Wayne, Limits of the learning curve, *Harvard Business Review*.