3. Descriptive results and corollaries

In this section we discuss the descriptive findings of the assessment of the performance of the coal-based power plants. The assessment considers all plants that were operational in the 30-month period from September 2017 to February 2020. Nearly 194 GW of capacity is considered in the assessment. Following the descriptive assessment, we present the parametric representation of station heat rate (SHR) and variable costs (VC), which will then be used in assessing physical and financial performance for the reassigned generation mix, in the subsequent section.

3.1 How are thermal power plants utilised?

Figure 1 illustrates the significant increase in power generation capacity. India has witnessed a huge capacity addition between 2010 and 2020. Nearly 65 per cent of the capacity as of March 2020 was installed in the previous ten-year period. We also note that 39 GW of capacity has been operating well past the economic life assumptions used in the determining the tariffs and returns on investment (CERC, 2014). Table 1 shows that half of the coal assets installed in the past decade have been done by the private sector. The remaining half of the installed capacity was equally shared between the central and state governments.

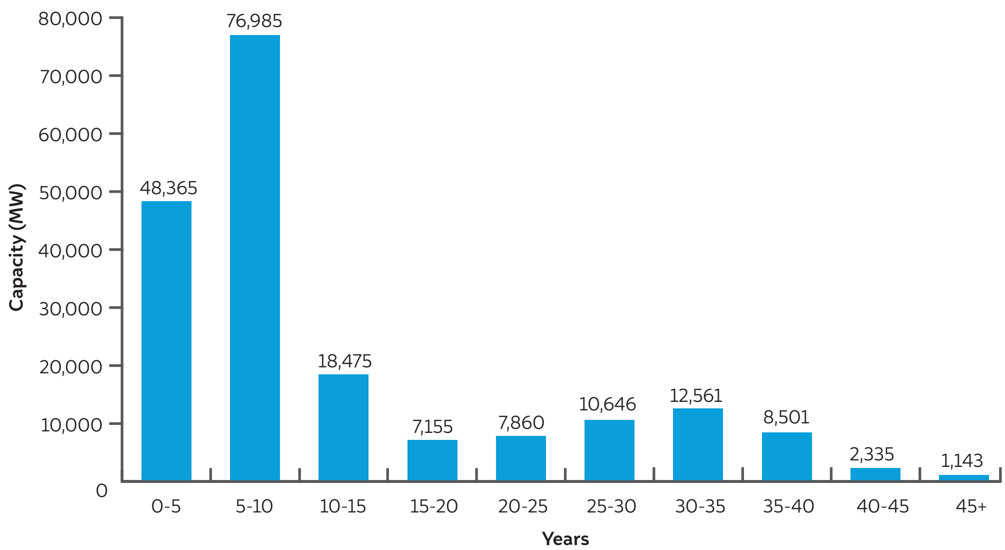

Figure 1
More than 125 GW of coal-based generation has been commissioned in the last ten years

Source: Authors' analysis from CEA daily generation reports

Table 1 Private sector investments have been the major driver of capacity addition in the last decade

Age group	Central sector (%)	Private sector (%)	State sector (%)	Total (%)
0–5 years	8	9	8	25
5–10 years	8	24	8	40
10–15 years	3	2	4	10
15–20 years	2	0.1	1	4
20–25 years	1	0.4	2	4
25–30 years	2	0.3	3	5
30–35 years	4	0.1	3	6
35–40 years	1	0.3	3	4
40–45 years	0	0.1	1	1
45+ years	0.2	0	0.4	1
Total	30	36	34	100

Source: Authors' analysis of CEA monthly installed capacity reports

It would be expected that newer plants being more efficient should be generating a higher share of the electricity in the system (as compared to their capacity share), as there would be economic gains from efficient generation. However, Figure 2 shows that the plants less than 10 years of age contribute a lower share to the total generation (62 per cent) than to installed capacity (65 per cent). Concomitantly, older plants contribute a disproportionately larger share of the generation as is clear from the illustration. As we show, it is contracting and other factors that determine this, and not efficiency.

Figure 2 Generation share of young plants are lesser than their respective capacity shares

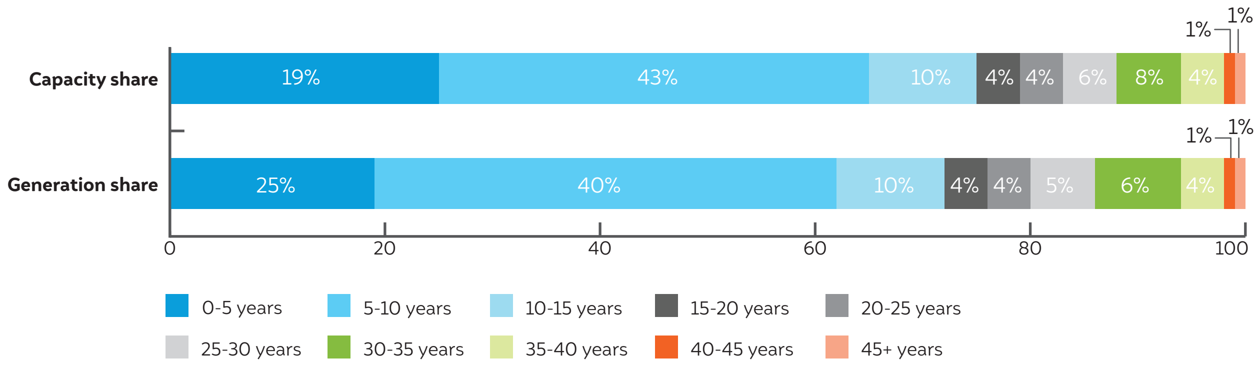

Source: Authors' analysis from CEA daily generation reports

In order to understand which segment of the plants are utilised poorly, we calculated a weighted average PLF for each category. Figure 3 indicates that the new plants have had very low PLF in the range of 40–60 per cent, while the older plants had high PLFs ranging between 75 and 85 per cent. On the whole, the state plants have been utilized less across age categories. The overall utilisation of the thermal fleet during the assessment period stood at a low of 58.5 per cent. This is clearly much lower than the envisioned PLFs for profitable operations of thermal assets.

Figure 3 State-owned plants show consistent under-utilisation across age groups

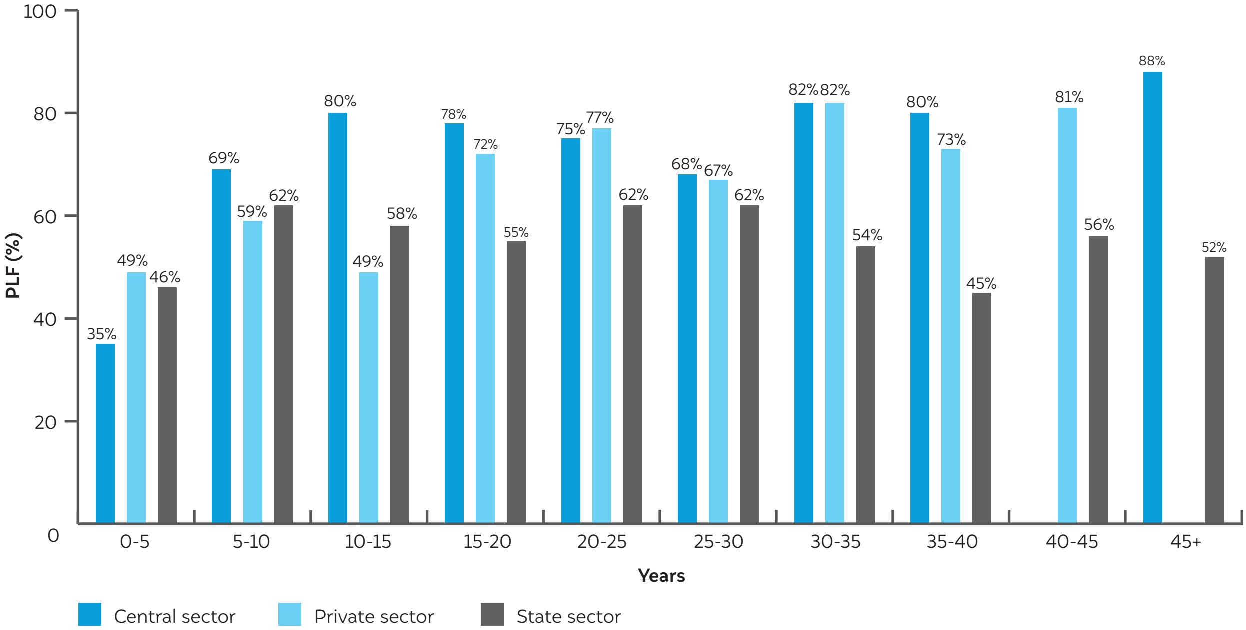

Source: Authors' analysis from CEA daily generation reports[3]

3.2 How efficient is the generating fleet of thermal power plants?

There is no consistent recording of data on the efficiency of thermal power plants at a high temporal resolution. What is available is an aggregate annualised metric, reported in tariff petitions filed by power plants and in an irregular CEA publication (with a lag of two years or more) that goes by the name 'Annual Thermal Performance Review'. Neither of these are useful to actually arrive at determinants of efficiency as many factors change over the course of a year and cannot be seen in aggregate. We set out to gather data on SHR at a higher temporal resolution. We accomplished it primarily by superimposing monthly coal consumption with the coal quality delivered, and then converting it to an efficiency metric by accounting for the electricity generated in each month.

The overall efficiency of the thermal operating fleet over the 30-month period stood at a paltry 29.7 per cent as per our calculations. The corresponding SHR was 2,898 kcal/kWh. This is particularly worrying as the improvement in the aggregate heat rate of the fleet over the years has not been commensurate with the pace of improvement in technology. This is not to say that plants did not operate more efficiently at all. A total of 29 plants exhibited an overall efficiency of more than 37 per cent (an SHR lower than 2,300 kcal/kWh) across many months in the analysis period. The median age of these plants was just a little over five years. Clearly, there are plants that are capable of performing more efficiently if the operations and circumstances allow them to.

The parametric estimation of determinants of SHR (kcal/kWh) was done through a panel regression with the independent variables being plant characteristics such as average unit size (MW), the average age of units (years), ownership (state or private, with central as the base), plant PLF (%), and share of imported coal in supply (%). We explored other variables such as measures of variability in (daily) plant loading, as theory suggests that deviation from base-load operation decreases plant efficiency. The daily variation

3 No plants in the central-owned/40–45 years category and private-owned/45+ years category is currently functioning. Also, the total capacity in 15–20, 20–25, 25–30, 30–35, 35–40 and 40–45 age groups of private sector and 45+ group of state and central sectors are insignificant (less than 1 GW).

in loading did not seem to have any relationship with SHR. We present the results of the panel regression below.

shr_kcalkwh	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
avg_unitcapacity	-.9257423*	.1175411	-7.88	0.000	-1.156119	-.6953658
plf	-330.5712*	24.58023	-13.45	0.000	-378.7476	-282.3948
age_baseline	7.640537*	2.240974	3.41	0.001	3.248308	12.03277
ownership_code						
Private Sector	-23.88326	63.81838	-0.37	0.708	-148.965	101.1985
State Sector	-73.21475	61.39452	-1.19	0.233	-193.5458	47.1163
import_share	225.4556*	92.20298	2.45	0.014	44.74112	406.1701
_cons	3417.812	80.99265	42.20	0.000	3259.069	3576.555

Figure 4
Age, unit capacity, and PLF are key determinants of SHR

Source: Author's Analysis

*Note: * indicates significant at 95 per cent confidence level*

The assessment in Figure 4 suggests that younger plants and large unit sizes have a beneficial impact on SHR. A detailed plot on how SHR varies with age can be seen in Figure A1 of the annexure. Ceteris paribus, a 660 MW unit, in comparison to a 300 MW unit, will have a heat rate lower by 300 kcal/kWh. Given that super-critical units are also identified beyond a capacity threshold, they are alone not useful in explaining the variation and correlated with the average capacity metric. Similarly, a 10-year-old plant will have a heat rate that is 75 kcal/kWh lower than a 20-year-old plant. An improvement in PLF by 20 per cent (in absolute terms) improves the heat rate by 65 kcal/kWh. Clearly, the most significant impact is made by unit size, and newer vintage plants are of an increasingly higher size, as would be expected with technology development.

3.3 Categorisation of the variable cost of coal plants

While the common perception is that older plants are cheaper because their fixed costs are paid for (discussed later in Section 4.2), what we see is that older plants are also often cheaper on a variable cost basis (Figure 5). On account of their lower generation efficiency, it would be expected that their cost of generation at the margin would be higher, but that is not the case. The variability, indicated in Figure 5, states that the age does not have clear implications for variable cost, given other factors at play.

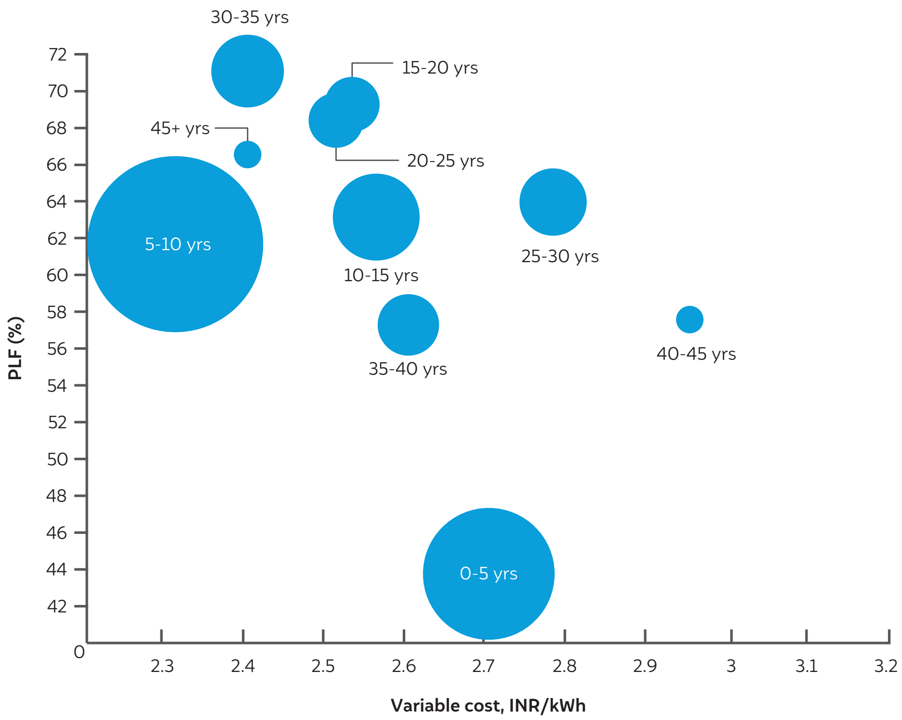

Figure 5
Despite having low variable cost, the PLF of 5- to 10-year group is low

Source: Authors' adaptation from CEA daily generation reports and Merit Order Despatch of Electricity for Rejuvenation of Income and Transparency (MERIT) state-wise daily summary data

Note: The bubble size represents the capacity share of each age group

Even in situations where younger plants, in line with efficiency arguments, have lower variable costs, we find that they are utilised to a lesser extent. How do we explain this? Discoms schedule power from the contracted generators based on the merit order dispatch (MoD) stack. Theory dictates that the generator with the lowest variable cost is dispatched first followed by the next lowest, and the iterative process continues until the energy or power demand is met, subject to technical constraints such as ramp rates and network capacity. The low PLF of certain plants (or vintages) should then reflect their (high) variable costs, indicating they are dispatched to a lesser extent. Interestingly, we see (Figure 5) that plants in the 5- to 10-year age group, which account for 40 per cent of the capacity share, despite having the lowest variable cost, have a lower PLF, compared to plants in the 20- to 35-year group. Similarly, plants in the 0–5 years bucket, despite having a lower variable cost than some of the oldest plants, operated at plant load that was 20 per cent lower. These two observations can be explained by the fact that plants that are contracted (either entirely or partially) are dispatched *only* to the extent they are contracted, as per the MoD. The uncontracted capacity either is typically treated as merchant power and sold on the exchange or through other mechanisms. They contribute to just about 10 per cent of the total procurement of electricity in the country (CERC, 2020). A significant share of capacity of newer plants remains to be contracted and, as a result, they don't get dispatched often.

3.3.1 What determines variable cost of generation at a plant?

On investigating the relationship between the variable cost and delivered cost of coal[4] (Figure 6), we found a high correlation between the two parameters. As expected, to a large extent, the delivered price of energy (INR/Mcal) is what determines the variable cost of electricity. This delivered cost of coal is largely a function of whether plants source coal from nearby mines or far away mines. Rail freight accounts for between 30 and 40 per cent of the delivered cost of coal (Kamboj & Tongia, 2018).

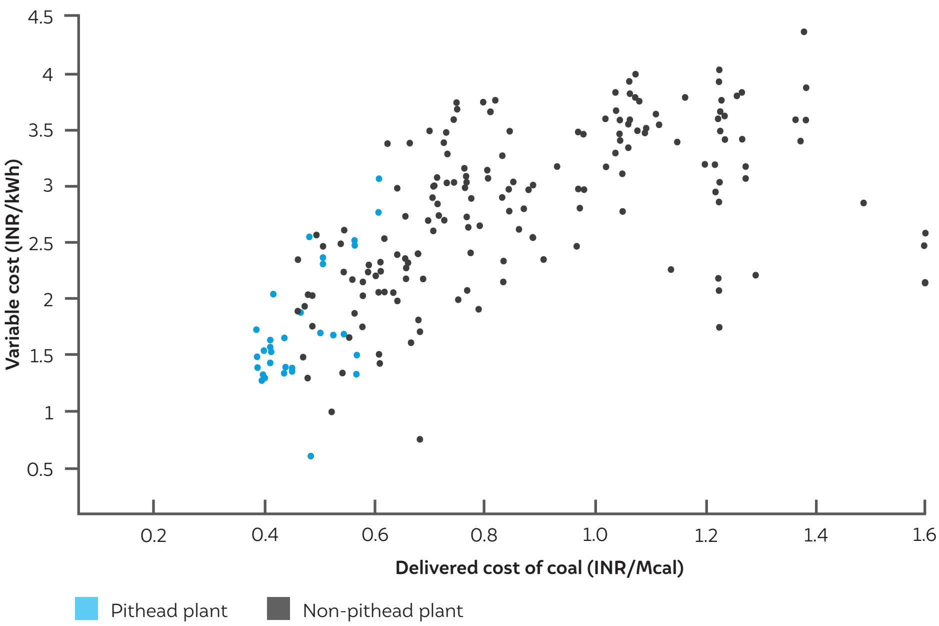

Figure 6
Variable cost of generation is driven to a large extent by delivered cost of coal

Source: Authors' analysis based on compilations from various generation tariff orders

However, in order to understand the determinants of variable cost (INR/kWh) more clearly, we carried out a linear regression-based assessment by considering independent variables such as auxiliary consumption (%), SHR (kcal/kWh), delivered coal price (INR/Mcal), average unit size (MW), and plant age (years). The results of the regression are presented in

4 Delivered cost of coal = Coal price (INR/ kg) / Gross calorific value of coal (kcal/kg).

Figure 7. We carried out this assessment for all the plants that have contracted capacity in part or full. Merchant generators are not considered in this calculation.

wtd_vc	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
shr_kcalkwh	.0001025*	.0000282	3.64	0.000	.0000473	.0001578
del_energy_cost	1.890785*	.0334197	56.58	0.000	1.825259	1.956312
age_baseline	.0074045*	.0012114	6.11	0.000	.0050293	.0097797
avg_unitcapacity	-.0006019*	.0000934	-6.44	0.000	-.0007851	-.0004188
aux_cons	2.827044*	.8023347	3.52	0.000	1.253896	4.400191
_cons	.6517435	.1273899	5.12	0.000	.4019686	.9015184

Figure 7
Other than coal price, age, unit size, SHR and auxiliary consumption were significant determinants of tariff

Source: Author's analysis

*Note: * represents variables that are significant at the 95 per cent confidence level.*

We observed that coal price has the most significant impact on the variable cost of electricity. The average delivered coal prices were at INR 0.74/Mcal. A 0.1 INR/Mcal decrease in coal price would reduce the variable cost by 0.19 INR/kWh, which is a substantial change. A 1 per cent decrease in auxiliary consumption would lower the variable cost by 0.03 INR/kWh. Similarly, replacing a 200 MW unit by 500 MW unit in the energy mix would cut down the variable cost by 0.12 INR/kWh. Running young and efficient plants too reduces the variable costs.

Given the significant upside associated with low-cost coal, we investigated the distribution of coal prices for various vintages of plants (Figure 8). We find that the plants of youngest vintage face a significant burden of high-priced coal and, as explained earlier, given their low levels of contracting, they face a double whammy of not being desirable even in an exchange or in merchant mode.

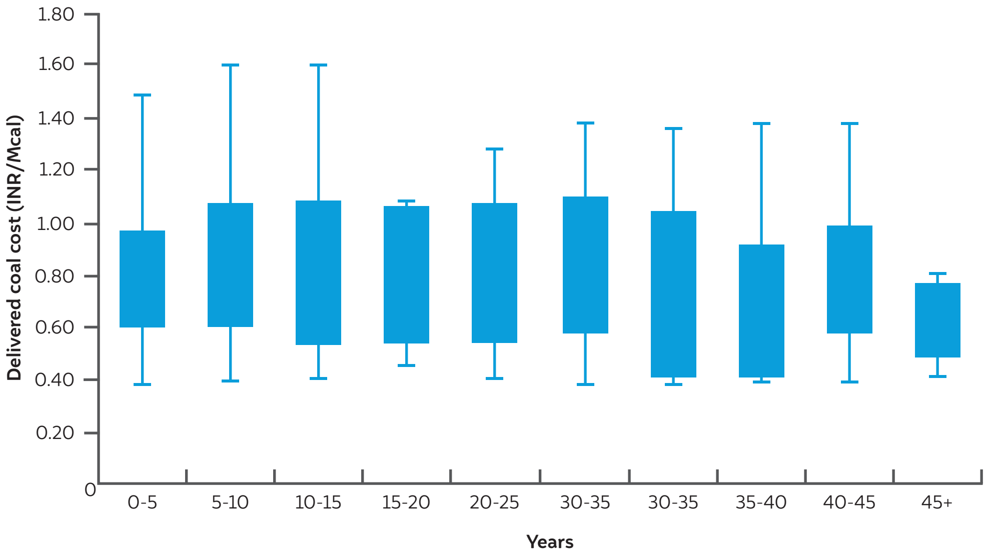

Figure 8
Plants between 30 and 40 years of age access the cheapest coal

Source: Author's analysis based on multiple tariff orders of generation companies

3.4 How does auxiliary consumption vary with vintage?

Another important implication of vintage in operational outcomes is the useful amount of energy that is actually available for commercial sale. Older plants have a higher auxiliary consumption, that is, power consumption within the plant itself (Figure 9). The same amount of energy, to be sold by a newer plant, would effectively require lesser generation to be undertaken, as more of the generated power will be available for sale. This in turn will reduce the amount of coal to be fired to generate electricity to that extent.