Out-of-merit costs and blackouts: Evidence from the Indian electricity market

Fiona Burlig, Akshaya Jha, and Louis Preonas*

November 2018

Preliminary and Incomplete
Not for citation or circulation

Abstract

In the United States, demand for electricity among utilities in the wholesale spot market is assumed to be perfectly inelastic. Consumers therefore face power outages only as a result of infrastructure failure – never because a utility does not purchase enough electricity to satisfy demand. This also implies that inefficiencies on the generation side of the market which raise price do not impact quantity consumed by retail customers. In this paper, we provide evidence that utilities participating in the Indian wholesale market are extremely price elastic: as prices rise, they purchase less power on the wholesale market, meaning that load shedding increases. Using data on plant-specific marginal costs, we document substantial deviations from first-best electricity generation, half of which can be explained by plant outages. These inefficiencies increase the wholesale price, and therefore contribute substantially to rampant blackouts.

Keywords: Electricity supply; India; market operations
JEL Codes: O18, Q41, L94

*Burlig: Department of Economics and Energy Policy Institute, University of Chicago. Saieh Hall for Economics #360, 5757 S. University Avenue, Chicago, IL 60637. Email: burlig@uchicago.edu. Jha: H. John Heinz III College, Carnegie Mellon University. 4800 Forbes Ave, Pittsburgh, PA 15213. Email: akshayaj@andrew.cmu.edu. Preonas: Department of Agricultural and Resource Economics and Energy Institute at Haas, UC Berkeley. 207 Giannini Hall, UC Berkeley, Berkeley, CA 94720. Email: lpreonas@berkeley.edu. We thank Steve Cicala, Michael Greenstone, Ryan Kellogg, Koichiro Ito, Nick Ryan, Catherine Wolfram, and seminar participants at the University of Chicago, Columbia – SIPA, Camp Resources, and the Heartland Environmental and Resource Economics Workshop for helpful comments and suggestions. Erin Kelley provided invaluable data acquisition support, and Jessica Jiang and Xiner Xu served as excellent research assistants. Burlig was generously supported by the National Science Foundation's Graduate Research Fellowship Program under grant DGE-1106400. Preonas was generously supported by Resources For the Future's Joseph L. Fisher Dissertation Fellowship. All remaining errors are our own.

1 Introduction

Electricity is an essential input to modern economic activity, and energy consumption is highly correlated with GDP in the global cross-section. Electricity demand in the developing world is projected to rise dramatically over the coming decades, as households move out of poverty and purchase electric appliances (Wolfram, Shelef, and Gertler (2012); Gertler et al. (2016)). Governments and development agencies invest billions of dollars annually to expand access to cheap, reliable electricity in low- and middle-income countries (Burlig and Preonas (2016)). Despite these investments, power outages remain a central feature of developing country electricity supply.

Blackouts are economically costly: firms cite unreliable power among their most important obstacles to growth. Power outages have been shown to cause reductions in manufacturing total factor productivity (Allcott, Collard-Wexler, and O'Connell (2016)) and increases in production costs (Fisher-Vanden, Mansur, and Wang (2015)). In the West, regulators place a premium on reliability; systems are dramatically overbuilt to avoid the potential for demand spikes to cause blackouts, and utilities are required to buy enough power to supply all of their consumers' demand, regardless of wholesale prices.

In many developing countries, outages are the result of a mismatch between demand and supply. While the United States has 3.3 kW of generating capacity per person, most countries in the developing world have per-capita generation capacity less than 1, with many countries having less than 0.1 kW per person. Even if every power plant in many of these countries is running full-bore, there is simply not enough generation capacity to supply demand without extremely high prices, which do not exist thanks to regulatory constraints. As a result, rolling blackouts are the norm.

In India, the setting of this paper, however, load shedding – the "politically correct" term for these blackouts – is common, despite a much more robust supply side than exists in much of sub-Saharan Africa. While there was a 1.7-fold increase in generation capacity between 2012 and 2018, consumption only grew by 1.4 times. In 2018-19, the government expects a 4.6 (2.5) percent overall (peak) energy surplus, but load shedding remains rampant.

Furthermore, despite the potential for shortages, over one quarter of generation capacity is declared to be under outage at any given time.

In this paper, we document two key features of the Indian electricity system which contribute to load shedding, despite a relatively large (for the developing world) generation sector. First, we show that electric utilities' demand for power is elastic: as wholesale electricity prices rise, utilities buy less, and instead clear the market on quantity, imposing blackouts. This is in stark contrast to a large existing literature on Western electricity markets, where demand is (assumed to be) perfectly inelastic (see Borenstein (2002) for an overview). As a result, anything which causes the price of electricity to rise, including market power, inefficient dispatch, and generating unit outages, will directly lead to blackouts. These blackouts cause allocative inefficiency, as consumers with a high value of power are unable to signal this to the utility.

Second, we document the existance of substantial out-of-merit generation costs on the supply side of the Indian electricity market. That is, we use novel data on electricity production and marginal costs of generation for each utility-scale power plant in the country between 2013 and 2017 to compare the total costs of generation with those that would have occurred under a least-cost counterfactual. That is, we "dispatch" power plants in order of lowest to highest cost, until there is enough aggregate supply to meet aggregate demand and clear the market. We compute the total short-run cost under this "least-cost" counterfactual, and we also calculate the total observed short-run costs of electricity generation based on each plant's actual production. By comparing the costs calculated under least-cost versus observed dispatch scenarios, we are able to quantify the cost gap between factual and counterfactual electricity supply. We find that this cost gap is large: even after accounting for interregional transmission constraints, and peak-versus-off-peak generation, there is a $1.4 billion, or 13.6 percent difference between least-cost and factual generation.

Generating unit outages are an important component of the difference between least-cost and observed dispatch. We find that the cost gap gap shrinks substantially – to 5.1 percent – when we take generator outages as given. However, two facts suggest that these outages may not be optimal from a social planner's perspective. First, while Indian plants do appear to need more maintanence than their American equivalents (Chan, Cropper, and

Malik (2014)), 27 percent of generation capacity is under outage on any given day – as compared with only 6 percent in the Texas electricity market. This extremely high outage incidence suggests there may be room for gains. Secondly, we find that these outages are uncorrelated with demand, suggesting that generators do not face incentives to produce power during high-value times.

The facts that generation costs appear to be substantially higher than first-best and that generating-unit outages are rampant are particularly important in light of the fact that utilities' demand for power is elastic. Any increase in cost over the first-best benchmark will raise the price of power in the wholesale market. This will in turn directly translate into less quantity being purchased, and more rolling blackouts. In ongoing work, we are performing a series of counterfactual exercises to quantify the impact of reductions in generating unit outages and decreases in out-of-merit costs on the quantity of electricity purchased by electric utilities, which has a direct relationship with load shedding.

This paper makes three main contributions to the energy economics literature. First, we document the existence of downward-sloping demand for electricity in a wholesale power market, a feature which is not present in any Western electricity market (Borenstein (2002)). Second, we add to a growing literature on electricity supply in developing countries.[1] We extend previous work on the impacts of power outages (Allcott, Collard-Wexler, and O'Connell (2016) and Abeberese (Forthcoming)) to study the determinants of outages. Finally, there is little existing work on the efficiency of electricity markets in developing countries (Jamasb, Nepal, and Timilsina (2015); Ryan (2017)). We quantify the impact of supply-side inefficiencies on purchased quantity in an important developing-country market.

This paper proceeds as follows. Section 2 describes the Indian electricity market. Section 3 presents our data. Section 4 outlines our empirical approach to (1) estimate the price elasticity of demand in the Indian wholesale market and (2) quantify the out-of-merit costs of generation in India, and Section 5 displays the results of this exercise. In Section 6, we discuss the importance of market power and generating-unit outages in driving a wedge between first-best and observed outcomes. Section 7 concludes.

1. There is a rapidly growing literature on electricity in the developing world. See for example Lee, Miguel, and Wolfram (2016), Burlig and Preonas (2016), and Dinkelman (2011) on rural electrification, among many others.

2 The Indian electricity market

In this section, we provide an overview of the major features of the Indian electricity market.

Physical structure

Fuel The majority of electricity generation in India comes from coal-fired generation sources. As of 2017, coal makes up 58.4 percent of this capacity, gas makes up 7.6 percent of this capacity, hydro makes up 13.5 percent of this capacity, nuclear makes up 2.1 percent of this capacity, and renewables makes up 18.2 percent of this capacity. Most coal- and gas-fired power plants procure input fuels via long-term Fuel Supply Agreements (FSAs), which are governed by fuel price regulations. Typically, fossil fuel plants are required to sign an FSA before they are allowed to begin operations. The majority of coal-fired power plants operate via linkages with Coal India, Ltd., the government-owned coal supply monopoly.[2] Recent shortages in coal and natural gas supply have been anecdotally linked with rolling blackouts and intermittent power supply.

Transmission As with any electricity market, transmission plays an important role in the Indian electricity sector. India is split into five electricity transmission grids, with inter-regional transmission lines between these five regions. The total interregional transmission capacity has increased rapidly in recent years, from 14,050 MW in 2007 to 75,050 MW in 2017. Expansions in transmission, including integrating all of India's transmission regions into one grid under the One Nation, One Grid scheme has been a major goal of the current government.

Transmission capacity is allocated by the National Load Despatch Centre in conjunction with regional and state Load Despatch Centres, using an administrative process that prioritizes long-term contracts, medium-term contracts, and finally the short-term electricity market (see Ryan (2017) for more details). Importantly, by prioritizing the transmission rights of long-term contracts, this limits the ability of generators to arbitrage differences

2. After our sample period, the government introduced the "Scheme to Harness and Allocate Kolya (Coal) Transparently in India" (a.k.a. Shakti) policy allocates *new* coal contracts to generation units based on an auction mechanism; the first auction ran in September 2017.

between contract prices and short-term trading prices. In other words, a long-term contract serves as both a forward financial commitment between traders *and* a de facto commitment to a physical allocation of electricity generation Ryan (2017).

Market structure

Contracts 90 percent of electricity in India is sold via medium- or long-term contracts between generation companies and wholesale consumers (e.g. distribution companies). The Central Electricity Regulatory Commission (CERC) approves the electricity tariffs received by generators on these bilateral contracts; this approval process is based on each plant's fixed and variable costs. 5 percent of generation is traded via short-term (less than one year from delivery) bilateral contracts.

Day-ahead market The Indian supply sector also includes a day-ahead wholesale market for power. Between 3 and 5 percent of power is traded on the day-ahead market in one of two exchanges: the Indian Energy Exchange (IEX) and the Power India Exchange Limited (PXIL). The IEX is substantially larger than the PXIL, but in both markets, generating units submit supply bids for each 15 minute interval of the following day. Discoms and industrial firms submit demand bids, the market is cleared where supply and demand intersect, and all suppliers are paid the market-clearning price.[3] Remaining balancing takes place using a mechanism called "unscheduled interchange," (UI) with regulated prices and quantities. Suppliers and consumers face penalties for overuse of UI.

3 Data

For this project, we digitized and assembled a novel dataset on the Indian electricity supply sector. Our empirical analysis draws on four main types of data. First, we use detailed daily plant-level generation, outage, and operation data, reported for the majority of India's utility-scale power plants by the Central Electricity Authority (CEA). Next, we construct

3. When transmission constraints bind, the market operators use an algorithm to clear sub-markets. See Ryan (2017) for details.

our own measure of power plant operating costs based on plant-specific heat rates and fuel consumption published by the CEA (a plant's heat rate is the amount of thermal energy in kcal required to produce one MWh of electricity). We incorporate detailed data on fuel prices to convert heat rates into marginal costs. These include fuel prices from the Ministry of Coal and Ministry of Petroleum and Natural Gas, hand-geocoded plant coordinates, and geospatial data on India's coalfields from the USGS—which we combine to calculate the shortest distance between each coal-fired power plant and the nearest coalfield. Finally, we have data from the IEX on market-clearing prices and quantities.[4]

Daily generation data The Central Electricity Authority (CEA) monitors the operations of all utility-scale (greater than 50 MW) fossil, hydroelectric, and nuclear power plants in India.[5] This include centrally owned, state-owned, and privately owned facilities. We obtain the CEA's Daily Generation Reports from January 1, 2013 (the earliest date available) to March 23, 2017.[6] Each report enumerates operational capacity, scheduled generation, and actual generation for each plant in each day-of-sample. The 485 plants in the CEA data make up 291 GW of India's 330 GW of electric generating capacity, with an average total generation of approximately 2.9 TWh per day (i.e. 2,900 GWh).

[Figure 1 about here]

Panel A of Figure 1 shows daily total generation summed over the plants in our sample while Panel B breaks this generation down by source. The vast majority of generation in the CEA data comes from the 196 coal-fired power plants, which average 2.2 TWh per day. The remainder comes hydroelectric (193 plants, 336 GWh per day), natural gas (62 plants, 121 GWh per day), nuclear (7 plants, 90 GWh per day), lignite (9 plants, 75 GWh per day), and liquid fuel-based generation (18 plants, 4.8 GWh per day).

4. We are in the process of digitizing full demand and supply curves from the IEX and PXIL.

5. Wind and solar resources instead fall under the Ministry of Renewable Energy. To the best of our knowledge, there exists no publicly available daily generation data from renewable sources. Non-hydro renewables account for less than 20 percent of India's electricity generation capacity. These technologies are non dispatchable and have extremely low marginal cost.

6. We expect our final analysis period to range from January 1, 2013 to December 31, 2017.

Outage data For each power plant, the CEA reports the MW of capacity that are under outage on any given day. These data are separated into "planned" and "unplanned" outages. A large amount of capacity is under outage on a regular basis. Figure 2 plots the fraction of capacity in our sample that is under outage. On average, over a quarter of generation capacity – 27 percent – is under outage on a daily basis.

[Figure 2 about here]

Marginal costs We construct an estimate of marginal operating costs using a variety of data sources. The CEA provides more detailed information on plant operations in the annual *Review of Performance of Thermal Power Stations*. In particular, each annual *Review* contains data on operating heat rates, a standard measure of electric generator efficiency measured in kilocalories per kilowatt-hour, for a large subset of coal-fired power plants. We obtained the 1997–2009 data from Chan, Cropper, and Malik (2014), and digitized the 2012–2014 *Reviews*, the most recent available at the CEA.[7] Because our analysis spans 2013 to 2017, we assign each plant its most recent heat rate that we observe. There are just 16 plants appearing in the *Reviews* for which the most recent heat rate was reported prior to 2012; for these plants, we obtained more recent heat rate data from tariff petitions to the Central Electricity Regulatory Commission. Indian coal prices vary by coal grade (i.e. coal's gross calorific value in kcal per kilogram), and we compute plant-specific coal grades by combining CEA heat rates with monthly coal consumption data from the CEA's Monthly Coal Reports. Ultimately, we have heat rate and coal grade data for 84 coal-fired plants and 7 lignite-fired plants, representing approximately 50 and 80 percent of each fuel's respective generating capacity from the CEA's daily generation data.[8]

We require data on fuel prices to convert plant-specific heat rates in to plant-specific marginal costs. Transportation costs contribute a substantial share of the total costs of coal

7. We thank the authors for sharing these data. We obtained access to hard copies of the the 2012–2014 *Reviews* on a recent trip to India.

8. Centrally owned and state-owned plants face stricter reporting requirements than private power producers (Chan, Cropper, and Malik (2014)). Hence, we observe heat rates for a non-random sample of power plants. In our empirical approach, we simply exclude plants for which do not have data from the calculation of costs for both observed and least-cost scenarios; thus, to the extent that unobserved plants are dispatched in the wrong order, we are under-stating the extent of the cost gap between these two scenarios. We discuss this and other measurement error issues in Section 4.

paid by power plants. In order to estimate these transportation costs, we geo-coded every power plant in our CEA sample by hand using a combination of CEA data and Google Maps. We also obtained a GIS map of the major coalfields in India from the U.S. Geological Survey (USGS), and match each coal-fired power plant with the closest coalfield. Then, we estimate the shortest as-the-crow-flies distance between each plant and its closest coalfield. Since the majority of plants are either pit-head plants, or transport coal via railway, we use coal-specific freight rates from the Freight Operations Information System of the Ministry of Railways to compute approximate transportation costs per kilogram of coal for each plant.[9]

Fuel prices are strictly regulated in India; state-owned coal suppliers set coal prices by grade of coal, which must be approved by India's Ministry of Coal. We digitized pithead coal and lignite prices from the Ministry of Coal's annual *Coal Directory of India*, which contains coal-grade specific prices.[10] We calculate coal costs as pithead prices plus transportation costs as well as royalties and other taxes.

For natural gas-fired power plants, we use heat rate data from the CEA's monthly gas report. Each monthly report presents total generation and natural gas consumption. We follow the Ministry of Natural Gas and Petroleum in assuming that the calorific content of natural gas is 10,000 kCal per standard cubic meter.[11] These data enable us to construct heat rates for 58 of the 62 gas plants in our daily CEA sample. We digitized natural gas prices from the Ministry of Petroleum and Natural Gas' annual *Petroleum and Natural Gas Statistics*, which yields a fuel price time series similar to our annual times series of grade-specific coal prices.

Wholesale market data The IEX posts data on market-clearing prices and quantities for every fifteen minute period. We downloaded these data for the duration of our sample. The IEX divides the market into five major regions: north, northeast, east, west, and south. We observe these data for each of the five regions, and in aggregate. Across our sample, the

9. We are currently working to incorporate geospatial information on the rail network to calculate transportation costs *along the rail network* (as in Preonas (2017)). This will modestly inflate our rail costs relative to those computed using as-the-crow-flies distances.

10. Pithead (a.k.a. mine-mouth) prices are the commodity price of coal at (close to) the point of extraction.

11. Monthly Gas Reports are available for 2012 and 2016–2017. We assign each plant its average observed heat rate over this period.

average price was 2,976 Rs/MWh, and the average market-clearing volume was 3,934 MWh. Figure 3 plots the time series of these two variables.

[Figure 3 about here]

In addition to these data, we are working to digitize information on aggregate IEX supply and demand curves. These data are available on the IEX website in `.png` form, so converting them to usable data requires a multi-step process. Figure 4 provides an example of a representative image.

[Figure 4 about here]

4 Empirical approach

4.1 Price elasticity of demand

We begin by quantifying the price elasticity of demand for electricity in the Indian wholesale elctricity market. To do so, we have two main approaches. First, we use an instrumental variables approach, where we use forced power plant outages as a supply shifter, to estimate the relationship between IEX prices and total quantity generated on a given day. In ongoing work, we use the demand curve data from the IEX to measure the price elasticity directly.

In order to estimate the price elasticity of demand for electricity in the Indian market, we are interested in regressing (log) quantity on (log) prices. This is of course the textbook example of simulteneity bias, so we require an instrument which will shift supply but not demand. We use the fraction of capacity that is on forced outage as one such shifter. In particular, we estimate the following system of equations using two-stage least squares, separately for each of the five IEX regions:

$$\log(Q)_{md} = \beta \log(\hat{P})_{mdt} + \gamma_m + \delta_t + \varepsilon_{mdt}$$
$$\log(P)_{mdt} = \nu \text{Share under outage}_{md} + \gamma_m + \delta_t + \eta_{mdt}$$

where Q_{dt} is total quantity in the electricity market in month of sample m on date d, and P_{mdt} is the price in the IEX market in month of sample m on date d in interval t. We use total quantity on the left-hand side to guard against the possibility of plants substituting generation for a long-term contract with production in the IEX and vice versa. This also allows us to demonstrate the relationship between the IEx prices and the full market, rather than just the 3-5 percent of demand that is fulfilled with the IEX supply. Share under outage$_{md}$ is the fraction of generation capacity that is under a *forced* outage. We define this share as the fraction of *inframarginal* capacity under a forced outage by considering outages only among plants whose marginal cost of generation is below the average price for each month-of-year-by-interval pair in our data. This prevents us from counting forced outages among peaker plants, rarely called upon, as relevant outages for this exercise, which improves precision. γ_m is a month-of-sample fixed effect, and δ_t is a fifteen minute interval fixed effect. ε_{mdt} and η_{mdt} are idiosyncratic error terms. We estimate Newey-West standard errors with a 7 day lag.

In order for this estimation strategy to be valid, two conditions must hold. The first stage requires that the share of capacity under forced outages must be correlated with the IEX price. We can test this empirically. As a check on this, we expect outages to be positively correlated with the price. The exclusion restriction is that the share of capacity under forced outages impacts quantity demanded only through the IEX price, or that forced outages are a supply shifter that does not impact demand. Given that we are including month-of-sample and interval-of-day fixed effects, this seems like a reasonable assumption. We also find no evidence that forced outages are correlated with temperature, a classic demand shifter.

4.2 Out-of-merit costs

Our empirical analysis of out-of-merit consists of three main steps. First, we construct the marginal operating cost for each power plant in our sample. Second, we compute the costs of marginal-cost-based—or "least-cost"—electricity dispatch. Next, we compute the costs of observed dispatch, and compute the total short run costs of observed deviations from least-cost dispatch. Finally, we conservatively account for both interregional and intraregional

transmission constraints and peak versus off-peak demand, in order to isolate the remaining out-of-merit cost wedge.

4.3 Marginal costs

To quantify marginal costs, and eventually aggregate variable costs, we invoke several standard assumptions in electricity markets. We assume that power plant i may generate any feasible quantity of electricity at time t, such that its generation Q_{it} falls between 0 and its binding capacity constraint $\overline{Q}_i$.[12] We also assume that for any quantity of generation $Q_{it} \in [0, \overline{Q}_i]$, plant i faces constant marginal costs MC_{it}. This is a reasonable approximation for plants' true marginal costs, which are driven largely by (1) its costs of purchasing fuel inputs P_{it}^f (e.g., the price of coal inclusive of transportation costs, which is independent of Q_{it}); and (2) its heat rate HR_{it}, or the calorific content of fuel needed to produce 1 MWh of electricity.[13]

We can express plant i's marginal costs per MWh at time t as:

$$MC_{it} = P_{it}^f \cdot HR_{it} + VOM_{it} \tag{1}$$

VOM_{it} represents variable operations and maintenance costs, including labor and non-fuel materials, which are typically much smaller than fuel costs P_{it}^f (Fabrizio, Rose, and Wolfram (2007); Cicala (2017)).[14] Nuclear plants face relatively lower fuel costs and non-trivial costs of operations and maintenance, compared to conventional fossil fuel generators. Hydroelectric plants typically incur close to zero marginal costs but face complex dynamic operating constraints.[15]

12. We compute $\overline{Q}_i$ as the 98th percentile of plant i's observed generation in each year, following Davis and Hausman (2016), to ensure that our capacity constraints are within the space of feasible plant operations. As a robustness check, we also define $\overline{Q}_i$ as the 80th percentile of plant i's observed generation.

13. Heat rates are the inverse of thermal efficiency, and plants with lower heat rates will tend to have lower marginal costs, all else equal. Energy economists typically assume that heat rates are constant across a plant's feasible quantities of generation (e.g. Cicala (2017)).

14. In the U.S. or Europe, plants may face (implicit or explicit) taxes on the emissions of sulfur dioxide (SO_2), nitrous oxides (NO_x), or carbon dioxide (CO_2). However, during our sample, Indian power plants did not face emissions taxes, so we exclude them from our marginal cost measures here.

15. Large hydroelectric dams often optimize multiple objectives, only one of which relates to electricity production (e.g., other water uses, dam release limits). Nuclear, coal, and (to a lesser extent) natural gas plants also face less extreme dynamic operating constraints, which we ignore in our analysis (Mansur (2008);

4.4 Least-cost and observed dispatch

Armed with a measure of each plant's marginal cost, we compute the difference between observed and least-cost dispatch. In doing this, we hold total quantity demanded fixed. This allows us to compute only the cost of out-of-merit generation. We also begin by abstracting away from transmission and ramping constraints, before incorporating interregional and intraregional constraints, as well as peak and off-peak generation, in an extension to our main analysis.

We calculate the cost of least-cost dispatch by finding the cost-minimizing allocation of electricity generation that could potentially meet observed market demand (or "load"):

$$COST_t^{LC} = \min_{\mathbf{Q}_t^*} \sum_i MC_{it} \cdot Q_{it}^* \quad \text{s.t.} \quad \sum_i Q_{it}^* = \sum_i Q_{it} \ , \quad 0 \leq Q_{it}^* \leq \overline{Q}_i \ \forall i \tag{2}$$

Here, Q_{it} denotes *observed* generation for plant i in time t, while $\mathbf{Q}_t^*$ is an $[I \times 1]$ vector of *counterfactual* least-cost generation Q_{it}^* for each of the same I plants. The solution to this constrained minimization problem simply sorts power plants from lowest-to-highest MC_{it}, and then "dispatches" each plant at its full capacity (i.e. $Q_{it}^* = \overline{Q}_i$) until aggregate idealized generation is equal to aggregate observed generation. We can calculate $COST_t^{LC}$ at the daily level, for each day t in our dataset. We start by assuming a fully unconstrained transmission grid, such that the idealized allocation of generation $\mathbf{Q}_t^*$ is not restricted by individual plant locations. Below, we describe how we incorporate transmission constraints and ramping into our least-cost counterfactuals.

Equation (2) effectively calculates the area under an idealized electricity supply curve, up to a given aggregate quantity $\sum_i Q_{it}$. Each level of $\sum_i Q_{it}$ implies a unique cost-minimizing vector $\mathbf{Q}_t^*$, which simply ranks plants in order of lowest-to-highest marginal cost. This ranking, known as the "merit order", serves as an efficient allocative benchmark and implies a weakly monotonically upward-sloping supply curve for electricity generation. Deviations from this idealized merit order increase total variable costs of electricity generation.

Reguant (2014)). Run-of-the-river plants store little to no water and provide non-dispatchable electricity at (virtually) zero marginal cost.

Calculating the cost of observed generation is more straight-forward, as we simply multiply each plant's observed generation (Q_{it}) by its marginal costs:

$$COST_t^{OBS} = \sum_i MC_{it} \cdot Q_{it} \tag{3}$$

Figure 5 illustrates the intuition behind calculating costs for both least-cost and observed dispatch. Gray shaded areas show the total variable cost associated with each blue dispatch curve. Observed dispatch meets market demand (vertical dashed lines) at greater total variable costs than least-cost dispatch.

[Figure 5 about here]

4.5 Cost gaps

We calculate the difference between realized vs. idealized costs for each day in our sample as:

$$COST_t^{DIFF} = COST_t^{OBS} - COST_t^{LC} = \sum_i MC_{it} \cdot (Q_{it} - Q_{it}^*) \tag{4}$$

$COST_t^{DIFF}$ summarizes the economic cost of deviations from least-cost electricity supply, or the potential decrease in total variable cost for day t. Importantly, the set of plants $i = \{1, \ldots, I\}$ entering into each cost calculation is a subset of the full population of electricity generating plants. We omit plants for which we do not observe data on marginal costs, meaning that $\sum_i Q_{it}$ is less than full market-wide electricity load for day t. By excluding a subset of plants from $COST_t^{DIFF}$, we implicitly assume that unobserved plants are efficiently dispatched and offer no potential gains from reallocation (i.e. $Q_{jt}^* = Q_{jt}$ for any plant $j \notin \{1, \ldots, I\}$). We also omit hydroelectric plants from our preferred cost calculations, due to the complex nature of characterizing theoretically "optimal" hydroelectric generation.[16] Hence, our calculations represent a lower bound on deviations from least-cost dispatch, if there exist additional cost-reducing potential reallocations that we do not observe.

16. If we included hydro plants, then the solution to Equation (2) would assume maximum hydro generation in all time periods due to hydro's extremely low marginal costs. In reality, dams cannot generate at 100 percent capacity due to several dynamic constraints—most notably, a finite supply of water.

4.6 Transmission constraints

The above cost difference calculation assumes unconstrained electricity transmission, yet physical transmission constraints may limit the flow of power between regions of the electric grid. For example, consider two electricity market regions (A and B) separated by a single transmission line with a finite flow capacity, but with otherwise unconstrained intraregional transmission. The efficient allocation of generation will dispatch plants in order of lowest-to-highest cost (ignoring plant regions) *unless* the AB transmission constraint binds. However, if the constraint binds and limits the interregional flow of electricity, then an efficient allocation of generation may force relatively high cost plants to operate in region A (in order to satisfy region A's demand) while relatively low cost plants in region B sit idle. In reality, India has five major transmission regions: the Northern Region, Western Region, Eastern Region, Southern Region, and North Eastern Region. Figure 6 displays these regions, along with the locations of all major power plants.

[Figure 6 about here]

To quantify the extent to which transmission constraints might explain the total cost gap $COST_t^{DIFF}$, we compute the costs of least-cost dispatch for each of India's five major electricity market regions (indexed by r):

$$COST_{rt}^{LC} = \min_{\mathbf{Q}_{rt}^*} \sum_{i \in r} MC_{irt} \cdot Q_{irt}^* \quad \text{s.t.} \quad \sum_{i \in r} Q_{irt}^* = \sum_{i \in r} Q_{irt} \ , \quad 0 \leq Q_{irt}^* \leq \overline{Q}_{ir} \ \forall i \tag{5}$$

Then, we sum across regions to calculate the total cost of least-cost dispatch:

$$COST_t^{LC_R} = \sum_r COST_{rt}^{LC} \tag{6}$$

This allows us to calculate the cost difference between observed dispatch and least-cost dispatch under interregional autarky:

$$COST_t^{DIFF_R} = COST_t^{OBS} - COST_t^{LC_R} \tag{7}$$

Here, we make the extreme assumption that each region is a completely separate market, or that there is zero interregional transmission capacity. We repeat this exercise using the 13 subregions of the Indian electric grid, which also accounts for *intra*regional transmission constraints by assuming interregional *and* intraregional autarky. The resulting cost difference calculations represent a lower bound on the short-run misallocation wedge, as trade between and within regions should only lower the cost of the idealized least-cost benchmark.

4.7 Ramping

Most power plants cannot be turned on instantaneously. Instead, after being started up, plants have a "ramping period" during which they cannot be used to generate electricity. If we do not account for the fact that plants cannot instanteously be used to meet demand, we may be overstating out-of-merit costs by counterfactually "dispatching" unusable plants. In an extension to our main result, we attempt to account for this by dividing the day into peak and off-peak hours. We assume 4 peak hours and 18 off-peak hours. We then use data on the ratio of peak versus off-peak generation to apportion the total observed quantity generated each day into peak and off-peak periods. We recompute out-of-merit costs for peak and off-peak separately, and sum these costs to recover the daily total.

4.8 Measurement error

It is possible that our plant-specific marginal costs data are imperfect. Measurement error in marginal costs will impact our total cost calculations: if our data understate (overstate) the marginal cost of some plants in our sample, we will compute total costs that are too low (high). However, we are primarily interested in calculating the cost of *deviations* from a least-cost dispatch order. The relative cost *differences* between least-cost dispatch and observed dispatch will be largely insensitive to classical and non-classical measurement error in plant-specific marginal costs, for two reasons. First, we use the same cost numbers to calculate observed costs and least-cost costs, meaning that inaccuracies will impact both estimates similarly. Second, and more importantly, our calculations aggregate costs across many inframarginal plants up to a single marginal plant. For measurement error to sig-

nificantly impact our findings, it would need to systematically overstate (understate) the operating costs of plants with low (high) costs.

5 Results

5.1 Demand elasticities

Using forced generating-unit outages as an instrument, we estimate the price elasticity of demand for electricity. We begin by estimating the first stage regression of log IEX price on the share of inframarginal capacity that is under outage, for each region separately. We find strong evidence that forced outages are associated with higher IEX prices in the North, Northeast, East, and West regions, with first-stage coefficients ranging from 0.96 to 3.30, and statistically significant at the 1 percent level. In the South, we do not find a statistically significant relationship. We think this is likely due to the spatial arrangement of the grid: while the North, Northeast, East, and West frequently import power, the South almost always exports. This means that the unit setting the market-clearing price for the South is unlikely to be in the South, which breaks the correlation between own-region outages and own-region prices. For the other four regions, however, we find strong positive relationships between outages and price, with F statistics ranging from 74.5 to 432.6. Table 1 shows both the first- and second-stage results.

[Figure 1 about here]

Armed with a strong first stage, we then estimate aggregate demand elasticities. We reject perfectly inelastic demand at the 1 percent level in the North and East, 5 percent level in the Northeast, and 10 percent level in the West. In the North, our estimated elasticity is 0.21; in the Northeast, 0.08; in the East, 0.09, and in the West, 0.08. These estimates are likely to be lower bounds, since the IEX likely matters most in peak (rather than off-peak) periods. Thi suggests that demand for power among Indian utilities is not perfectly inelastic, as in the West, but rather elastic. In the IEX itself, we find even stronger evidence of elastic demand, as demonstrated in Figure 4.

5.2 Out-of-merit costs

Next, we quantify the extent to which out-of-merit costs exist in the Indian power market. We do this by constructing marginal costs, calculating costs under least-cost dispatch, and calculating costs under observed dispatch. We then account for transmission constraints and ramping.

5.2.1 Marginal costs

We begin by constructing the marginal-cost-based merit order for the Indian power sector. Using our measures of marginal cost, we simply rank plants from lowest to highest. Figure 7 shows the results of this exercise, which effectively solves the cost minimization problem in Equation (2). We exclude hydropower from our analysis because of the complexities of dealing with temporal opportunity costs. Our merit order is sharply differentiated by fuel type: then low-cost nuclear is cheapest, then coal (with some nuclear mixed in), and then a mix of gas and lignite. This is perhaps unsurprising, as coal is cheaper than gas, and lignite plants tend to be less efficient. Given that our cost estimate includes only heat rates and fuel prices, most of the variation in marginal costs comes from fuel prices alone (rather than a combination of fuel prices and variation in operations and maintenance costs). Our constructed marginal cost also tend to be significantly less than cost reported by the Ministry of Power. In our constructed cost sample, the average coal plant has a marginal cost of $15.6 per MWh, the average gas plant has a marginal cost of $36.5 per MWh, the average lignite plant's marginal cost is $29.3 per MWh, and the nuclear plants range in cost from $4.2 per MWh to $18.7 per MWh.

[Figure 7 about here]

Coal plants make up the bulk of the merit order. Gas plants are more costly.[17] These marginal costs are broadly consistent with other estimates of the cost of generation in the United States (Davis and Hausman (2016)) and in India (Ryan (2017)).

17. We exclude the single liquid-fuel-based power plant from Figure 7. This plant, Basin Bridge, has a marginal cost of $270 per MWh.

5.2.2 Least-cost dispatch

Armed with these marginal costs, we can follow the procedure outlined in Section 4 to compute generating costs of least-cost dispatch for each day in the sample. We first calculate total load on day t as $\sum_i Q_{it}$—that is, observed generation for each plant in our sample, for which we have nonmissing data. As described above, we exclude hydroelectric plants from these results, because they face complex dynamic constraints. By dropping hydro generation from the sample prior to constructing our measure of total load, we essentially assume that they are always dispatched optimally.

To compute the cost of least-cost dispatch, we "dispatch" plants in order of least cost, at full capacity, until we reach total observed generation. This creates a vector $\mathbf{Q}^*_{it}$ where for a given marginal cost MC^*_t:

$$\begin{aligned} Q^*_{it} &= \overline{Q}_i && \text{if } MC_{it} < MC^*_t \\ Q^*_{it} &\in \left[0, \overline{Q}_i\right] && \text{if } MC_{it} = MC^*_t \\ Q^*_{it} &= 0 && \text{if } MC_{it} > MC^*_t \\ \textstyle\sum_i Q^*_{it} &= \textstyle\sum_i Q_{it} \end{aligned}$$

This vector $\mathbf{Q}^*_{it}$ solves the constrained minimization problem in Equation (2), and allows us to calculate the cost of least-cost dispatch on day t as $\sum_i MC_{it} \times Q^*_{it}$.[18]

[Figure 8 about here]

Figure 8 displays the results of this exercise. Panel A shows the total costs of least-cost dispatch over the course of our sample. $23.3 million per day, with a standard deviation of $1.9 million. This translates to an average variable cost of approximately $14.8 per MWh, from dispatching in order of our constructed marginal costs.

Panel B displays the marginal cost of the marginal plant (or MC^*_t). As expected, the cost of the marginal plant tracks closely with the total cost, and with total daily load shown in Figure 1.[19]

18. Costs in our marginal cost data are reported in rupees. We convert to USD with an exchange rate of 60 rupees to 1 dollar.

19. Note that total daily load from Figure 1 is not identical to total daily observed generation in these calculations, as we only use plants for which we have marginal cost data (excluding hydroelectric generation) in the results shown in Figure 8.

5.2.3 Least-cost versus observed dispatch

Next, we compute actual costs following Equation (3). We simply multiply each plant's observed generation by its marginal cost, and sum over all plants for each day in the sample. This allows us to calculate the difference between the costs of observed and least-cost dispatch, as in Equation (4). Figure 9 presents these cost differences. Panel A shows a kernel density of the daily cost gap between observed and least-cost dispatch, and Panel B plots the daily time series of cost difference per MWh.

[Figure 9 about here]

We compute average *observed* costs of \$28.0 million per day, or \$17.6 per MWh. Combining these results, we find an average cost difference between observed and least-cost-dispatch-based costs of \$4.7 million per day (ranging from \$3.0–\$6.4). The cost gap between observed vs. least-cost dispatch represents approximately 16 percent of total observed costs. Annually, this amounts to a cost gap of \$1.7 billion per year. Put differently, the cost difference between observed dispatch and least-cost dispatch amounts to over 16 percent of realized variable costs. This suggests that there may be substantial short-run gains from reallocation in the Indian electricity market.

We calculate a substantial cost gap between observed and least-cost dispatch, indicating that electricity generation departs markedly from the cost-based merit order. This cost gap could reflect systematic departures, wherein many plants are frequently dispatched by a rule of thumb that is imperfectly correlated with marginal cost. On the other hand, a few large high-cost plants could regularly generate out of merit and be responsible for most of the deviations we observe.

For each day in the sample, we classify a plant as "out of merit" if the plant generates despite a least-cost generation of zero, or vice versa. In other words, we assign plant i to be out of merit on day t if $Q_{it} > 0$ and $Q^*_{it} = 0$ *or* if $Q_{it} = 0$ and $Q^*_{it} > 0$. Panel A of Figure 5 reports a histogram for the percent of days that plants generate out-of-merit, for all plants in our sample. This reveals that the majority of plants rarely generate out of merit, while a specific subset of plants frequently generate out of merit. In fact, roughly 25 plants generate out of merit on at least 90 percent of the days in our sample. This suggests that

reducing out-of-merit generation for these few plants could yield substantial short-run cost reductions.

[Figure 10 about here]

We then compute the level difference between observed generation and least-cost generation for each plant-day, or $Q_{it} - Q^*_{it}$. Panel B of Figure 10 plots each plant's difference between observed and least-cost generation against its marginal costs. Unsurprisingly, departures from least-cost generation are strongly correlated with marginal costs of generation: plants with higher marginal costs are more likely to generate out-of-merit. In particular, we find that while low-cost plants generate slightly too little on average, there are a few higher-cost plants that generate substantially more than our least-cost benchmark.

5.2.4 Transmission constraints

Whereas our results have thus far imposed a strong assumption of no transmission constraints, we now consider the opposite extreme by assuming autarky for both regions and subregions of the Indian electric grid. Using Equation (5), we calculate the costs of least-cost dispatch separately for plants in each of 5 regions (or 13 subregions), and aggregate the total cost differences across regions (or subregions). Figure 11 reports three distributions of daily cost gaps, assuming respectively: (i) unconstrained transmission (reproduced from Panel A of Figure 9); (ii) interregional autarky, with unconstrained transmission within subregions; (iii) interregional *and* intraregional autarky.

[Figure 11 about here]

Under interregional autarky, we find that the average cost difference between observed and least-cost dispatch is $3.8 million per day, or 13.9 percent. This is only slightly less than the average cost difference under fully unconstrained transmission. In other words, interregional transmission constraints appear to account for less than 20 percent of the cost difference between observed and least-cost generation allocations. We see a similar result if we rescale cost differences by total generation. This comparison is striking: even after accounting for interregional transmission constraints, widely considered to be the most

important physical limitation in electricity markets, the vast majority of misallocation in electricity generation remains.

Under both interregional *and* intraregional autarky, we calculate an average cost gap of 11 percent. This implies that transmissions constraints between regions/subregions explain less than 31 percent of the cost difference between observed and least-cost allocations. Having accounted for the physical constraints of electricity transmission both across and within regions, the remaining 11 percent cost gap represents a conservative estimate of the short-run out-of-merit cost wedge in Indian electricity supply.

5.3 Peak vs. off-peak generation

When we account for peak vs. off-peak generation in addition to inter-regional autarky, we find a cost gap of 13.6 percent, as compared with the 13.9 percent which just accounts for inter-regional transmission. This suggests that failing to account for ramping constraints is not substantially affecting our main results.

6 Discussion

In this paper, we argue that demand for electricity among electric utilities in India is relatively price elastic. This implies that any cost increases on the generation side of the power market will lead directly to load shedding. We find that the observed dispatch of power plants in India results in annual total variable costs that are $1.4 to $1.7 billion higher than the costs implied by least-cost dispatch. As a point of comparison, Cicala (2017) estimates that a shift from traditional command-and-control dispatch to market-based dispatch in the U.S. reduced electricity generating costs by $3 billion per year, in a much larger market. This suggests that there are important out-of-merit costs in the Indian market, which are contributing to load shedding. In this section, we explore the hypothesis that market power is driving our results and discuss the importance of generating-unit outages.

6.1 Market power

First, we consider the extent to which firms might be exercising market power, which is common in wholesale electricity markets even when supply is not heavily concentrated among a few firms (Borenstein (2002)). In the Indian context, Ryan (2017) estimates average quantity-weighted markups of roughly 20 percent above marginal costs, for electricity sold on the short-run day-ahead market during 2009–2010.[20] However, market power is unlikely to explain a significant share of the short-run cost gap, for two reasons. First, roughly 90 percent of electricity is sold on medium- and long-term contracts with fixed tariffs, which preclude the exercise of short-run market power. Second, recent transmission capacity expansions have likely decreased firms' abilities to exercise market power, by effectively increasing the number of plants able to supply electricity to a given location of demand (Borenstein, Bushnell, and Stoft (2000); Ryan (2017)).

Figure 12 plots deviations from least-cost generation against plants' marginal cost, splitting plants by vintage (Panel A), total capacity of their owning entities (Panel B), and ownership type (Panel C). If market power were a major driver of the cost gaps we calculate, we would expect to see plants with above-median firm capacity to withhold output from their relatively high-cost plants in order to drive up market prices and earn greater inframarginal rents form their relatively low-cost plants. However, Panel B reveals that out-of-merit generation for high-cost plants is comparable across large vs. small firms; if anything, larger firms appear to operate their low-cost (inframarginal) plants relatively *less* than small firms. This suggests that market power is unlikely to be driving the bulk of the cost gap between observed and least-cost dispatch.

[Figure 12 about here]

6.2 Generating-unit outages

Power plants must undergo maintenence, which causes them to be offline for lengthy periods. Up until this point, we ignore plant reports of outages when computing the cost gap between

20. In levels, Ryan (2017) estimates an average short-run price-cost markup of roughly \$8.5/MWh.

least-cost and observed dispatch. When we instead re-estimate least-cost dispatch while taking plant outages as given, we find a striking result: the cost gap reduces to approximately $1.3 million per day, or only 5.1 percent – less than half of our interregional autarky estimate. Given that there may be remaining measurement error causing our estimates of the cost gap to be upward-biased, such as local transmission constrants and grid stabilization requirements, we argue that a 5.1 percent gap is relatively small. This would seem to suggest that there is little room in the Indian power grid for cost improvements.

However, though Indian plants need to be taken offline for upgrades more often than their US counterparts (Chan, Cropper, and Malik (2014)), this alone is not nearly enough to explain the rampant outages we observe. We see that 27 percent of generating capacity is under outage on an average day — as compared with only 6 percent in Texas. Furthermore, these outages appear to be uncorrelated with demand: moving from a 25th to a 75th percentile demand day is only associated with less than a 4 percent change in the capacity that is under outage. This suggests that there is substantial potential for a reduction in the total number of outages or a change in their timing which could lead to lower costs of power supply.

7 Conclusion

Economic development is strongly linked to increased energy consumption. Load shedding may therefore hamper economic growth and impose undue burdens on poor consumers. In this paper, we demonstrate that Indian utilities have a sharply downward-sloping demand curve for electricity, meaning that as prices rise, they purchase less power, and instead impose rolling blackouts on their consumers. The fact that utilities are price-elastic means that anything which raises costs on the generation side of the market will directly lead to load shedding. We document evidence of substantially elevated costs in Indian electricity supply by comparing least-cost dispatch to observed dispatch, and find a cost gap of between 13 and 16 percent. This gap is cut in half when we take plant outages as given, but these outages are rampant and appear to be inefficient. Given that roughly 90 percent of electricity generation occurs on long-term contracts, allowing power plants to more flexibly arbitrage the day-

ahead spot market would likely alleviate a substantial share of the short-run misallocation we are finding.

In ongoing work, we are (1) generating a new dataset of wholesale market demand and supply curves, which will allow us to calculate the elasticity of demand for power in the wholesale market only; and (2) performing counterfactual exercises where we lower the cost of supply or reduce the frequency of generating-unit outages to estimate the increase in quantity supplied that could be achieved with reasonable improvements to market efficiency.

References

Abeberese, Ama Baafra. Forthcoming. "Electricity Cost and Ferm Performance: Evidence from India." *Review of Economics and Statistics.*

Allcott, Hunt, Allan Collard-Wexler, and Stephen D O'Connell. 2016. "How do electricity shortages affect industry? Evidence from India." *American Economic Review* 106 (3): 587–624.

Borenstein, Severin. 2002. "The Trouble with Electricity Markets: Understanding California's Restructuring Disaster." *Journal of Economic Perspectives* 16 (1): 191–211.

Borenstein, Severin, James Bushnell, and Steven Stoft. 2000. "The Competitiveness Effects of Transmission Capacity in a Deregulated Electricity Industry." *RAND Journal of Economics* 31 (2): 294–325.

Burlig, Fiona, and Louis Preonas. 2016. "Out of the Darkness and Into the Light? Development Impacts of Rural Electrification." Energy Institute at Haas Working Paper No. 268.

Chan, Hei Sing Ron, Maureen L. Cropper, and Kabir Malik. 2014. "Why Are Power Plants in India Less Efficient than Power Plants in the United States?" *American Economic Review* 104 (5): 586–590.

Cicala, Steve. 2017. "Imperfect Markets Versus Imperfect Regulation in U.S. Electricity Generation." National Bureau of Economic Research Working Paper No. 23053.

Davis, Lucas, and Catherine Hausman. 2016. "Market Impacts of a Nuclear Power Plant Closure." *American Economic Journal: Applied Economics* 8 (2): 92–122.

Dinkelman, Taryn. 2011. "The Effects of Rural Electrification on Employment: New Evidence from South Africa." *The American Economic Review* 101 (7): 3078–3108.

Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram. 2007. "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on U.S. Electric Generation Efficiency." *American Economic Review* 97 (4): 1250–1277.

Fisher-Vanden, Karen, Erin T. Mansur, and Qiong (Juliana) Wang. 2015. "Electricity Shortages and Firm Productivity: Evidence from China's Industrial Firms." *Journal of Development Economics* 114:172–188.

Gertler, Paul J., Orie Shelef, Catherine D. Wolfram, and Alan Fuchs. 2016. "The Demand for Energy-Using Assets among the World's Rising Middle Classes." *American Economic Review* 106 (6): 1366–1401.

Jamasb, Tooraj, Rabindra Nepal, and Govinda R. Timilsina. 2015. "A Quarter Century Effort Yet to Come of Age: A Survey of Power Sector Reforms in Developing Countries."

Lee, Kenneth, Edward Miguel, and Catherine Wolfram. 2016. "Experimental Evidence on the Demand for and Costs of Rural Electrification." National Bureau of Economic Research Working Paper No. 22292.

Mansur, Erin T. 2008. "Measuring Welfare in Restructured Electricity Markets." *Review of Economics and Statistics* 90 (2): 369–386.

Preonas, Louis. 2017. "Market Power in Coal Shipping and Implications for U.S. Climate Policy." Energy Institute at Haas Working Paper No 285.

Reguant, Mar. 2014. "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets." *Review of Economic Studies* 81 (4): 1708–1742.

Ryan, Nicholas. 2017. *The Competitive Effects of Transmission Infrastructure in the Indian Electricity Market.* Technical report. National Bureau of Economic Research Working Paper No. 23106.

Wolfram, Catherine, Orie Shelef, and Paul Gertler. 2012. "How Will Energy Demand Develop in the Developing World?" *Journal of Economic Perspectives* 26 (1): 119–138.

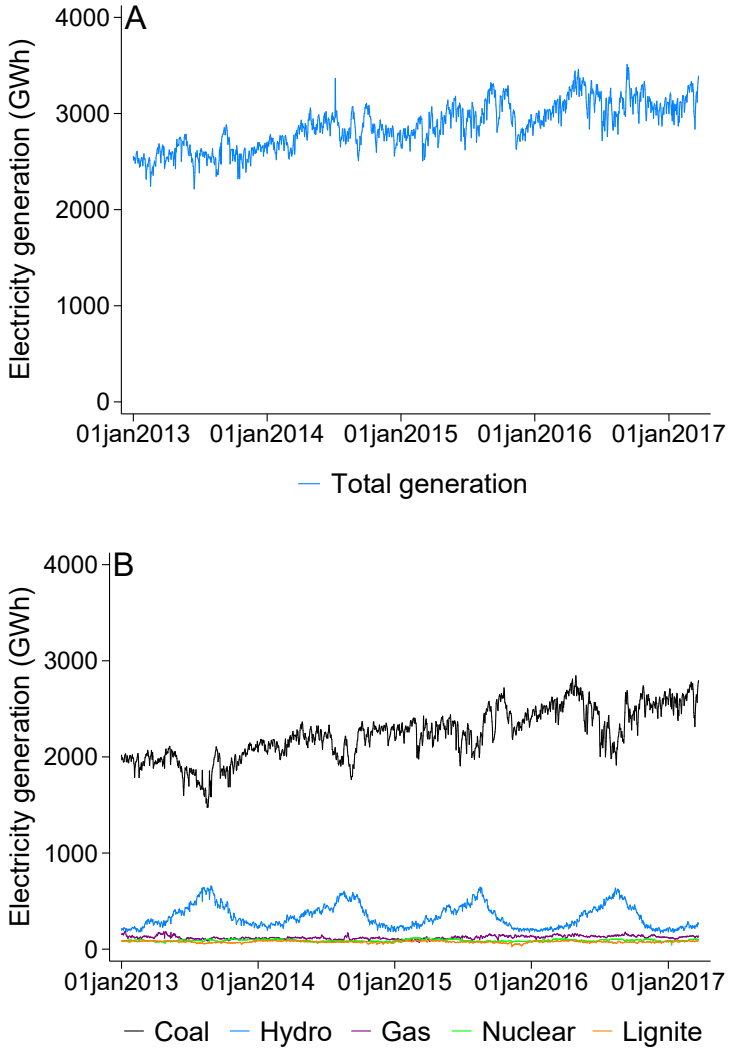
Tables and Figures

Table 1: Elasticity of power demand

	Aggregate $\log(Q_{\text{region}})$			
	North	Northeast	East	West
	(1)	(2)	(3)	(4)
IV: $\log(P^{\text{IEX}}_{\text{region}})$	-0.21^{***}	-0.08^{**}	-0.09^{***}	-0.08^{*}
	(0.07)	(0.04)	(0.03)	(0.04)
First stage:				
Share capacity on forced outage	0.96^{***}	3.30^{***}	0.77^{***}	1.38^{***}
	(0.11)	(0.16)	(0.08)	(0.13)
Month-of-sample FEs	Yes	Yes	Yes	Yes
15-minute interval FEs	Yes	Yes	Yes	Yes
Date-interval observations	146,592	146,096	146,304	146,880
First-stage F-stat (IV)	74.5	432.6	101.6	105.4

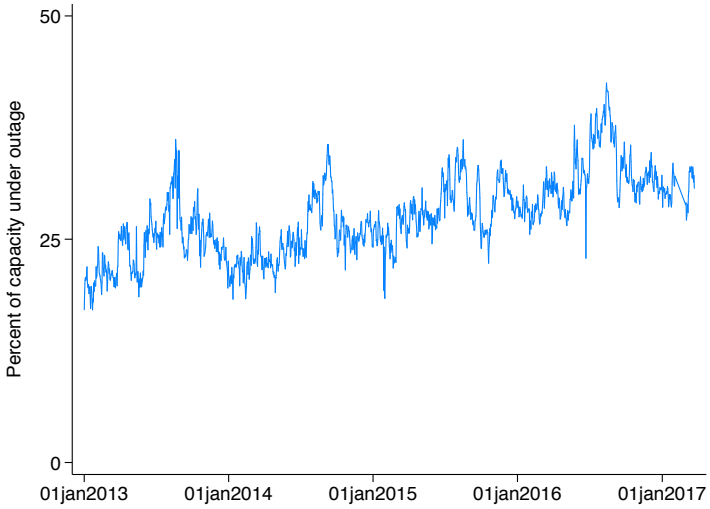
Notes: This table reports demand elasticities, generated by regressing the log of quantity demanded on log IEX price. We instrument IEX price with the share of capacity under forced outage. Each column shows a distinct region; regressions are estimated separately. Standard errors are Newey-West with a 7 day lag.

Figure 1: Electricity Generation in India – CEA Daily Data



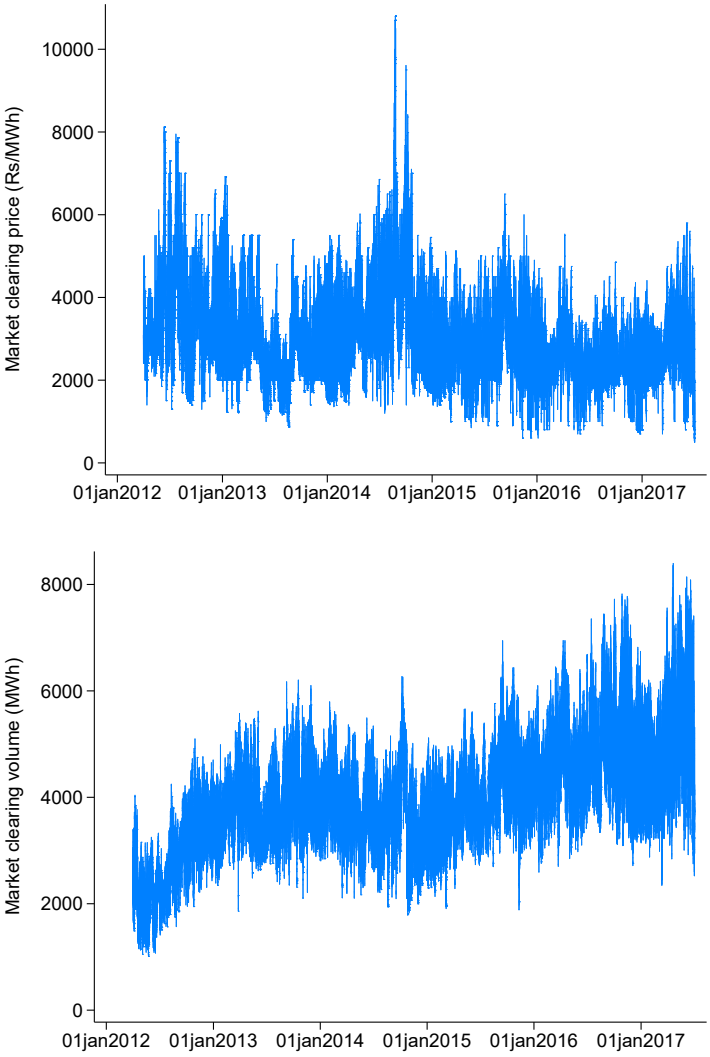
Notes: This figure displays data from the Central Electricity Authority's Daily Generation Reports. Panel A shows total electricity generation in GWh from January 1, 2013 through March 23, 2017. The 485 plants in these data average 2.9 TWh per day. Panel B breaks these plants out by fuel type. The 196 coal-fired power plants generate by far the most electricity, averaging 2.2 TWh per day. The remainder of generation comes from 193 hydroelectric plants (336 GWh per day, but highly seasonal); 62 natural gas plants (121 GWh per day); 7 nuclear plants (90 GWh per day); 9 lignite plants (75 GWh per day); and 18 liquid-fuel-based plants (4.8 GWh per day).

Figure 2: Capacity under outage – CEA Daily Data



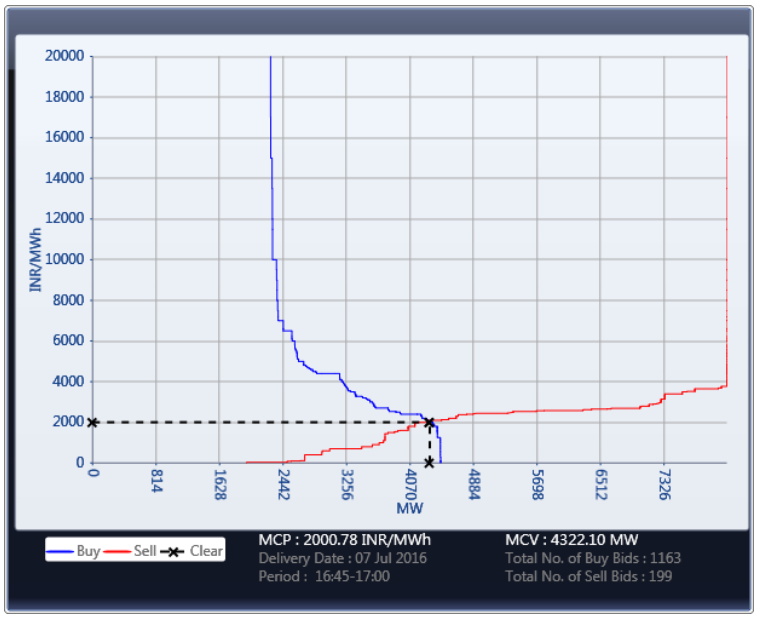
Notes: This figure displays data from the Central Electricity Authority's Daily Outage Reports. On average, 27 percent of capacity is under outage each day.

Figure 3: Wholesale market data – IEX



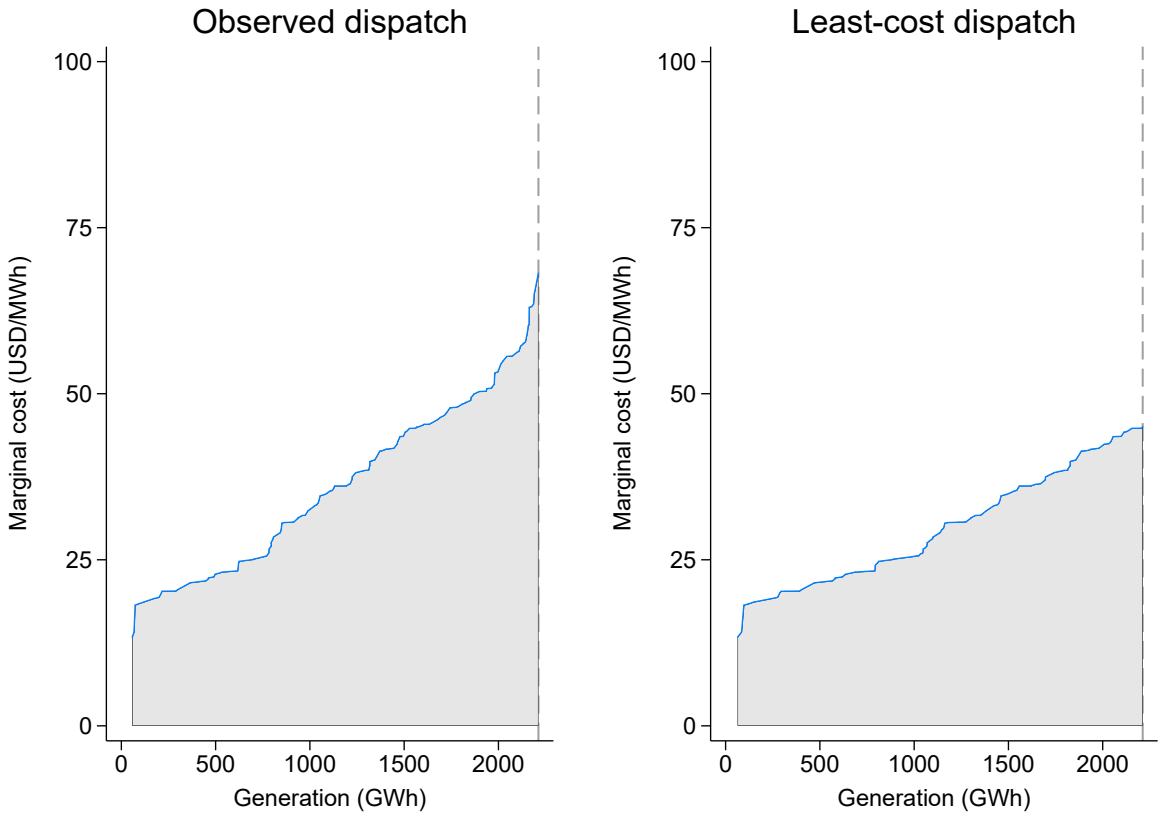
Notes: This figure displays data from the Indian Energy Exchange, India's largest wholesale electricity market. The top panel shows market clearing prices for each interval; the bottom shows the corresponding market clearing quantities.

Figure 4: Example IEX demand and supply curves



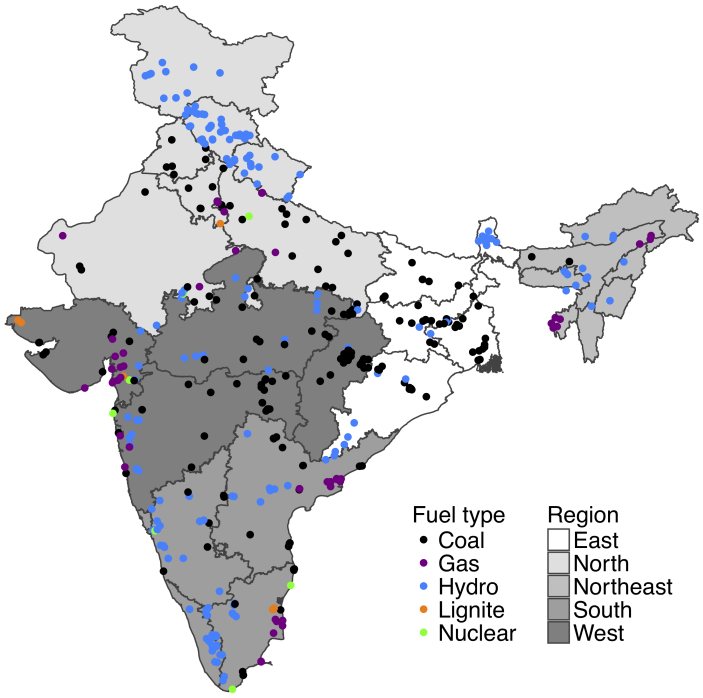
Notes: This figure displays an example day from the Indian Energy Exchange data.

Figure 5: Costs of Observed vs. Least-Cost Dispatch



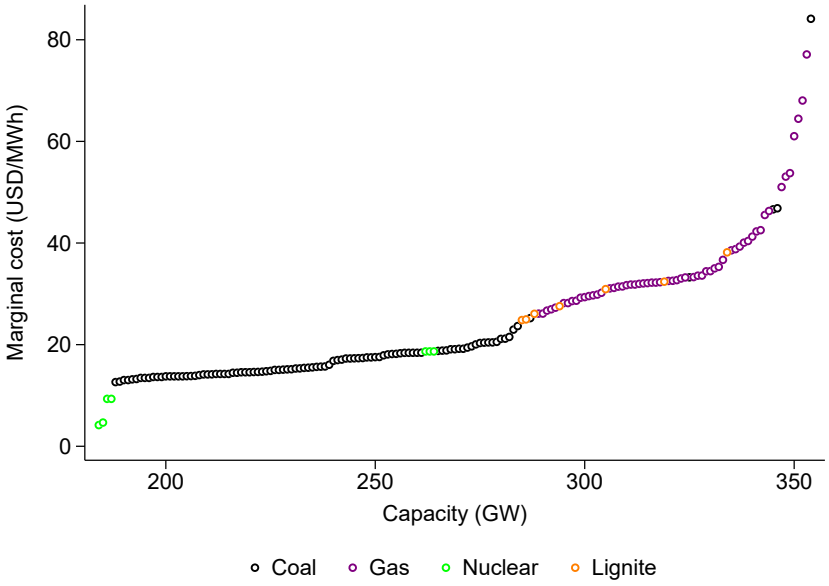
Notes: This figure illustrates the empirical approach underlying our main estimates using data from March 12, 2015. In the "Observed dispatch" panel, we take all the plants that were observed to be running on this date, and plot the quantity they generate in increasing order of marginal cost. The blue line shows the "supply curve" of plants, as ordered by marginal cost. The gray area is our estimate of total observed costs – in this case, $78.9 million. The dashed line shows the total load on this day, of 2.2 TWh. The "Least-cost dispatch" panel, by contrast, orders all plants from lowest to highest marginal cost, and dispatches plants to their capacity in this order until we reach 2.2 TWh of generation. Again, the gray area represents the total costs of meeting this load: $65.3 million, or $13.6 million lower than the cost of observed dispatch. It is clear from this figure that total costs under least-cost dispatch are lower than those under observed dispatch, and that this is because more low-cost plants are being dispatched to meet load under the least-cost-dispatch scenario.

Figure 6: Electricity Market Regions of India



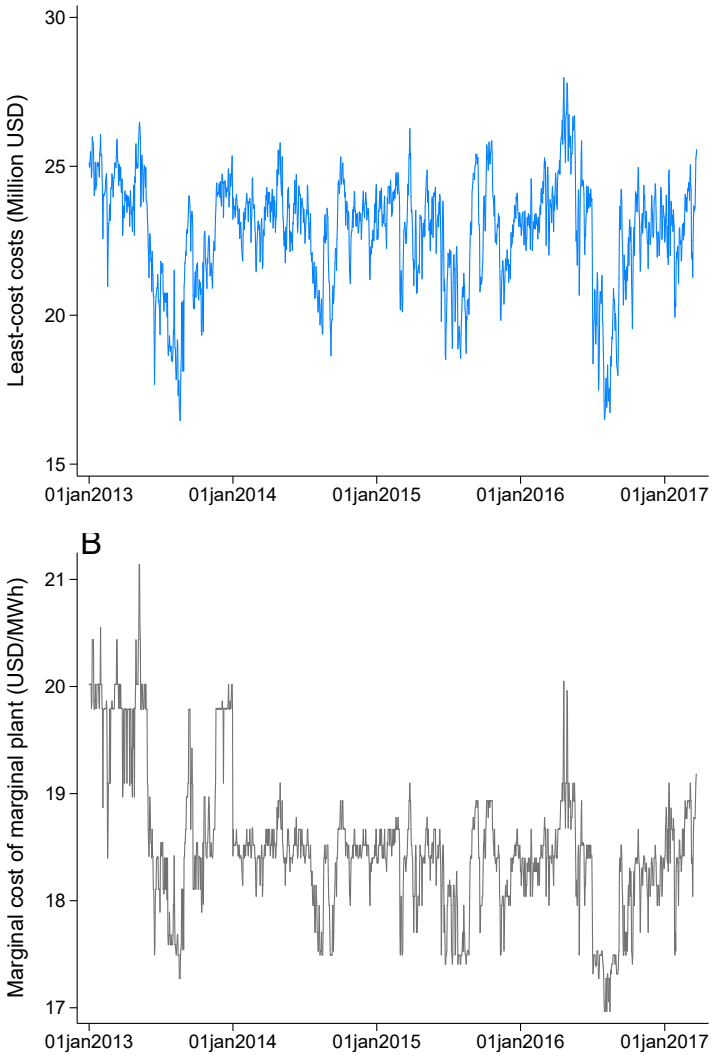
Notes: This figure displays the locations of the 485 power plants that appear in the CEA's Daily Generation Reports. The plant colors indicate the fuel type – coal, gas, hydroelectric, lignite, and nuclear. States are shaded by their major transmission region – North, North East, East, West, and South.

Figure 7: Marginal-cost-based merit order



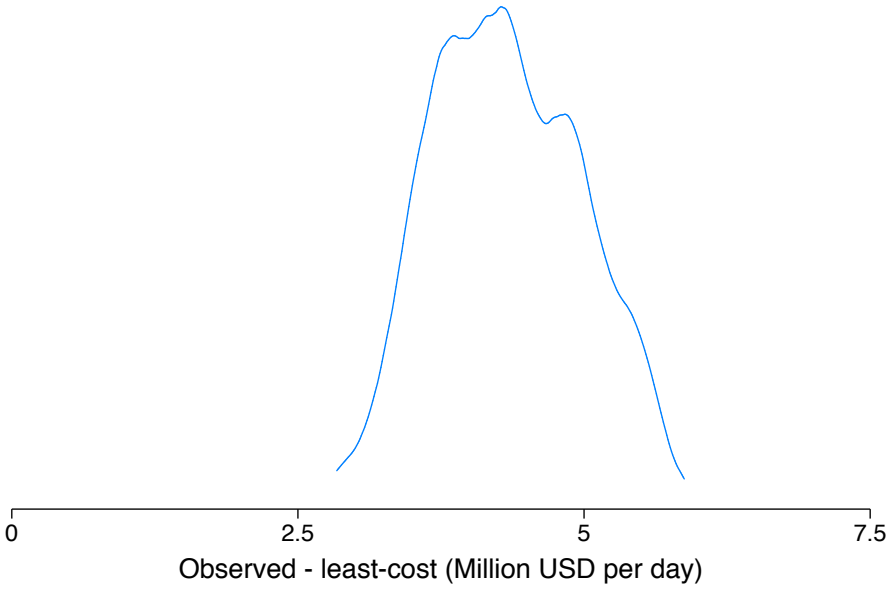
Notes: This figure shows the Indian electricity supply merit order, where we rank plants by their marginal cost according to our own constructed marginal cost estimates. This figure shows all fuel types together, with colors denoting fuel types. Each dot represents 1 GW of capacity; larger plants will be represented by multiple dots. India has a large amount of hydroelectric generating capacity - though this is hamstrung in reality by dynamic considerations, which we assume has a marginal cost of zero. The next entries in the merit order are low-cost nuclear facilities; followed by coal-fired plants, interspersed with several other nuclear power plants; more coal; and a mix of coal-, natural gas-, and lignite-based plants.

Figure 8: Costs of least-cost dispatch



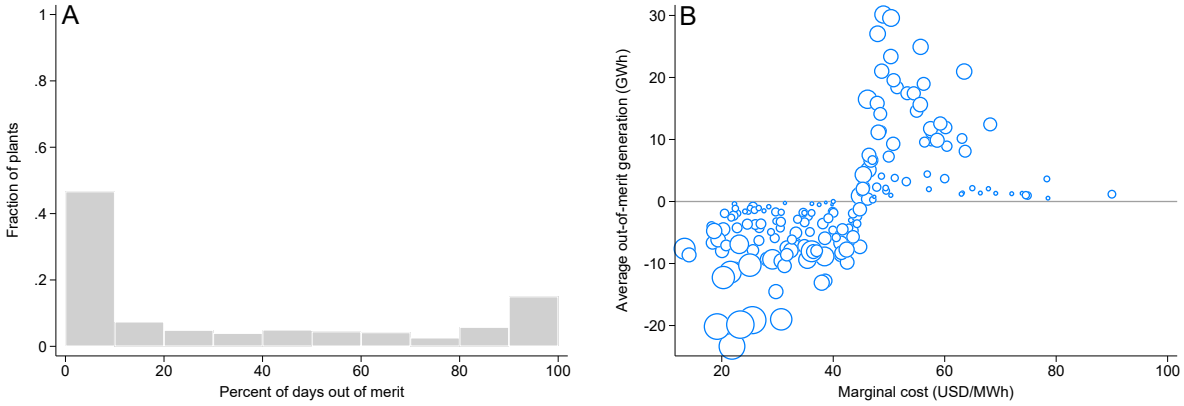
Notes: This figure displays our estimates of the cost of least-cost generation in India. As described in Sections 4 and 5, we construct the least-cost-dispatch-based cost estimates by ranking plants from lowest to highest marginal costs, "dispatching" each plant to generate up to its capacity, until we meet total daily load. Panel A displays the total short-run costs of least-cost generation for the plants in our sample. Panel B shows the marginal cost of the marginal plant under least-cost dispatch. As expected, higher total cost days (which correspond to higher overall load days, as shown in Figure 1) must rely on higher-cost plants to be operating.

Figure 9: Difference between observed and least-cost-dispatch-based costs



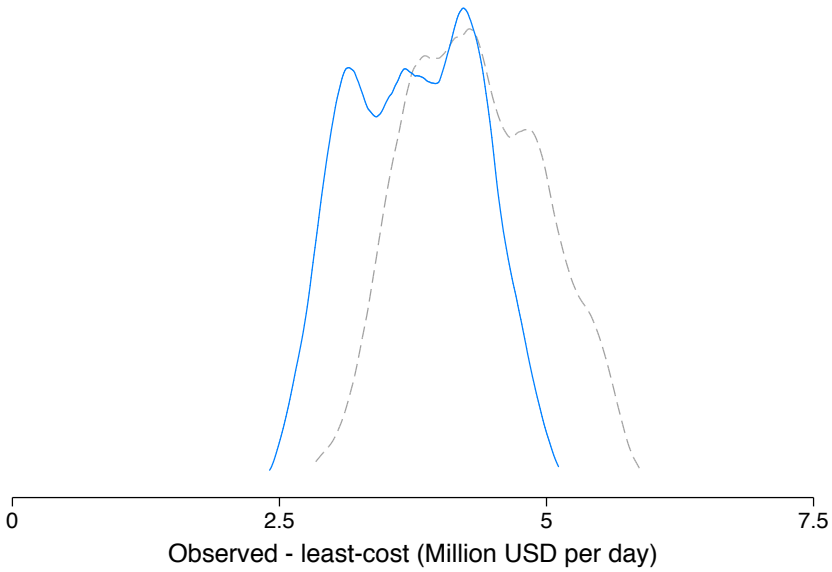
Notes: This figure shows the differences in cost between least-cost and observed dispatch in the Indian electricity market. Panel A shows the distribution of the cost difference between least-cost dispatch and observed dispatch across the sample. Panel B shows the cost gap *per MWh* over the sample.

Figure 10: Out-of-merit generation



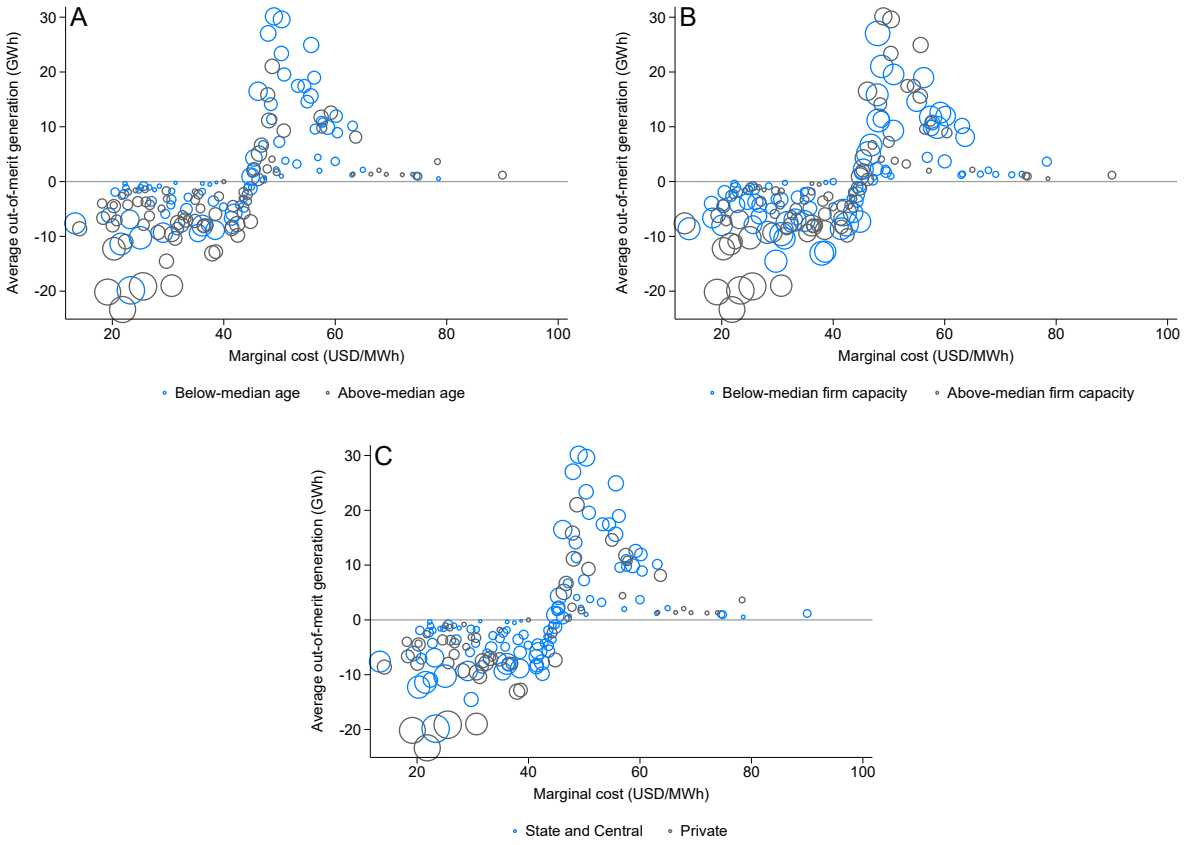
Notes: This figure explores out-of-merit generation in least-cost dispatch as compared with observed dispatch. Panel A shows the percent of days that each plant generates out-of-merit, or is observed to be operating in the CEA data when it should not be operating according to our least-cost dispatch order. While most plants rarely generate out-of-merit, around 20 percent of plants generate in the observed dispatch order but not in the least-cost dispatch order more than 90 percent of the time. Panel B shows the average difference between observed generation and generation under least-cost dispatch as compared with marginal costs per MWh. There is a strong correlation between marginal costs and out of merit generation: low-cost plants appear to be generating too little, while higher-cost plants are generating too much. Both panels are weighted by capacity.

Figure 11: Role of inter-regional transmission constraints



Notes: This figure displays the cost difference (in millions of dollars per day) between least-cost dispatch and observed dispatch under two different scenarios. The dashed gray line shows the distribution of cost differences when we create the least-cost dispatch order by ranking *all* plants by their marginal cost. This essentially assumes that electricity from any plant can meet load anywhere in the system. By contrast, the solid blue line displays cost differences from an exercise in which we treat India's five transmission regions – as shown in Figure 6 – as five closed economies. That is, all demand in a given region can be met by plants in that region only. This estimates an upper bound on the impact of inter-regional cost differences on the total cost gap between least-cost and observed dispatch, under the extreme assumption that there is no transmission capacity connecting regions. Even under this strong assumption, we find that less than 20 percent of the cost gap is driven by inter-regional transmission constraints, suggesting that more than 80 percent of the difference in cost between least-cost and observed dispatch remains unaccounted for.

Figure 12: Out-of-merit generation – Heterogeneity analysis



Notes: This figure shows the relationship between out-of-merit generation and marginal costs in dollars per MWh, weighed by capacity (as in Figure 10). Here, we explore heterogeneity in out-of-merit generation. Panel A shows heterogeneity by plant age. Blue dots show plants below the median age, built after 2007; and gray dots show older plants, built in or before 2007. We use the age of the average unit in the plant to define plant age. We see no systematic difference in the relationship between marginal cost and out-of-merit generation between older and newer plants. Panel B demonstrates a preliminary test for market power: are (high-cost) plants owned by larger firms (gray dots) more likely to generate out of merit than plants owned by smaller firms (blue dots)? We again see no discernible patterns. Finally, Panel C shows examines out-of-merit generation by firm type: government run (central and state, blue dots) versus privately owned (gray dots). Ownership does not appear to predict out-of-merit generation. The same is true if we split the state and central sector into their own categories.