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

Fiona Burlig, Akshaya Jha, and Louis Preonas*

March 2019

Preliminary and Incomplete
Not for citation or circulation

Abstract

Rolling blackouts are a common feature of electricity markets in developing countries. If distribution utilities are price sensitive, providing less electricity when the marginal costs of generation are high—that is, if wholesale electricity demand is elastic—, then supply-side distortions that raise wholesale prices will reduce the quantity of electricity that ultimately reaches consumers. In this paper, we document two key facts about the Indian electricity market. First, demand for electricity in the wholesale market is downward sloping and substantially more elastic than in Western markets. Second, we demonstrate that power plant outages are an important contributor to out-of-merit generation. We show that the timing of these outages in the full market, where 90 percent of electricity is provided via long-term contracts, is unresponsive to demand shifters. However, plants participating on the spot market increase their quantity in response to these same shifters, suggesting that incentives can play an important role in outage decisions. We conduct a simple back-of-the-envelope exercise in which we add capacity under outage into the spot market supply curve, to simulate a shift to market incentives, and find that relatively small changes in outages among low-marginal-cost plants can lead to decreases in the wholesale price and increases in quantity supplied. These changes in quantities would be enough to reduce Indian utilities' claimed shortages by 15 percent, suggesting that there may be room for market mechanisms to improve outcomes in the Indian wholesale power sector.

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

*Burlig: Harris School of Public Policy and Energy Policy Institute, University of Chicago. 1307 E 60th St, 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 Economics and Energy Policy Institute, University of Chicago. 5757 S. University Avenue, Chicago, IL 60637. Email: lpreonas@uchicago.edu. We thank (without implicating) Severin Borenstein, Steve Cicala, Michael Greenstone, Koichiro Ito, Ryan Kellogg, Shaun McRae, Nick Ryan, Catherine Wolfram, and seminar participants at the University of Chicago, Columbia – SIPA, Camp Resources, the Heartland Environmental and Resource Economics Workshop, SETI at Duke, the Indian Statistical Institute, and the Energy Institute at Haas' Energy Camp for helpful comments and suggestions. We benefited from conversations with Rajesh Mediratta and Shilpy Dewan of the IEX. 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. These outages are economically costly; for example, 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 many developing countries, blackouts occur because electricity demand exceeds the maximum aggregate production capacity available to meet this demand. However, in our empirical setting of India, blackouts are common despite there being ample production capacity to meet aggregate electricity demand. This is in part because distribution utilities must sell electricity to consumers at subsidized prices that can be below the true marginal cost of supply. Rather than sell electricity at a loss, these utilities will often choose to leave some demand unmet.[1]

This is in stark contrast to electricity markets in the West, where regulated utilities are required to provide power to all consumers and cannot impose discretionary blackouts (Borenstein (2002)). Since wholesale electricity demand is close to perfectly inelastic in Western markets, supply-side distortions such as market power can increase electricity *prices*, but do not alter the *quantity* of electricity ultimately delivered to consumers.[2] This intuition does not carry over to the developing world. In developing countries where utilities exhibit downward-sloping electricity demand, any supply-side distortions that increase the

1. Compounding this problem are low bill payment rates, where utilities only receive payment at these subsidized prices from a fraction of consumers.

2. Previous work has documented the exercise of market power in electricity markets in California (Borenstein, Bushnell, and Stoft (2000), Borenstein, Bushnell, and Wolak (2002), and Wolak (2003, 2007)), the U.S. Midwest (Mercadal (n.d.)), Texas (Hortacsu and Puller (2008)), Great Britain (Wolfram (1999)), and Spain (Reguant (2014) and Ito and Reguant (2016)).

marginal costs of electricity generation will increase electricity prices *and* decrease electricity quantities—thereby increasing the rate of blackouts faced by consumers. To the extent that demand-induced blackouts ration electricity for consumers who are willing to pay, these blackouts create allocative inefficiency in retail electricity provision.

This paper quantifies the extent to India's demand-induced blackouts are caused by one major supply-side distortion: power plant outages. Indian power plants have notoriously high outage rates (Chan, Cropper, and Malik (2014)), with roughly 26 percent of total capacity unavailable to generate on any given day. Excessive plant outages shift electricity supply inward, which raises prices; since utilities have downward-sloping wholesale electricity demand, these higher prices reduce the quantity of power that reaches consumers. Our analysis proceeds in three steps. First, we document downward-sloping demand in India's day-ahead wholesale electricity exchange, while also estimating how aggregate electricity demand responds to changes in the day-ahead market clearing price. Second, we identify plant outages as a prominent driver of high variable costs of electricity generation in India, and show descriptive evidence that both the level and timing of plant outages are not correlated with economic factors or market conditions. Third, we examine counterfactual scenarios with fewer plant outages, to see the extent to which mitigating this supply-side distortion could potentially reduce blackouts and increase retail electricity provision.

First, we document that the aggregate demand curves submitted to India's largest electricity exchange are downward-sloping. We use 15-minute interval data on aggregate supply and demand curves submitted to the Indian Energy Exchange (IEX), a day-ahead wholesale electricity market that contributes roughly 3–5 percent of electricity generation in India. For all 96 15-minute intervals, the average wholesale demand elasticity at the market clearing price is between -1 and -2. As these elasticities only represent a small share of total wholesale power demand, we also estimate the elasticity of *aggregate* demand with respect to the day-ahead IEX price. Instrumenting for price with the share of capacity on unplanned outage (a plausible supply shifter), we find downward-sloping aggregate demand in 4 of India's 5 electricity supply regions, with elasticities ranging between -0.14 and -0.20. Hence, utilities not only purchase less electricity on the day-ahead market when prices are high—they also purchase less *total* electricity when day-ahead prices are high.

Second, we calculate the total variable costs of electricity generation, and compare these costs to counterfactual "least-cost" dispatch. In this "least-cost" scenario, we re-dispatch power plants in order of lowest-to-highest cost, allowing us to quantify the potential reductions in variable costs under this idealized (albeit unrealistic) counterfactual. We find an average daily cost difference of 16.0 percent between observed variable costs vs. "least-cost" variable costs, and this cost difference shrinks slightly to 13.6 percent when we account for interregional transmission and separate peak vs. off-peak demand. However, when we re-dispatch plants taking declared outages as given (i.e. removing all capacity under outage on each day), the cost difference shrinks to 5.1 percent. This suggests that lowering the rate of outages has the potential to substantially reduce the costs of Indian electricity generation.

We also provide evidence that the incentives power plants face matter for their outage decisions. In the full market (where 90 percent of power is traded on long-term contracts), we find that neither the frequency nor the timing of plant outages is correlated with economic factors, such as marginal costs, plant ownership sector, or temperature. This suggests that India's high outage rate is indeed *distorting* the supply-side of the electricity sector by raising the variable costs of generation. In contrast, among plants participating in the day-ahead market, we find evidence that plants respond to temperature by increasing supply offers. This suggests that plant incentives may be an important driver of outages in India.

Finally, we perform several back-of-the-envelope calculations to quantify the increase in electricity production that may result from the interaction between a reduction in supply-side outages and downward-sloping demand for electricity in the wholesale power market. To do this, we identify marginal plants under outage, and expand the IEX supply curve by adding in a fraction of the "unavailable" capacity. We then compute the new market clearing price and quantity by intersecting this expanded IEX supply curve with the demand curve. We find that adding a relatively modest 20 percent of capacity under outage at extremely low marginal cost plants (50 percent or less of the IEX market clearing price) reduces the IEX market clearing price by 200 rupees per MWh, and increases equilibrium quantity in the IEX by nearly 500 MW on average, or approximately 15 percent of the average peak shortage. This suggests that there may be room for improved incentives on the generation side of the wholesale electricity market to translate into an expanded power sector, through lowering prices faced by price-sensitive wholesale market buyers.

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.[3] We extend previous work on the impacts of power outages (Allcott, Collard-Wexler, and O'Connell (2016) and Abeberese (2017)) 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, and highlight the importance of generator incentives in driving these inefficiencies.

This paper proceeds as follows. Section 2 describes the Indian electricity market and presents our data. Section 3 shows that Indian wholesale electricity demand is downward-sloping. Section 4 illustrates how plant outages are likely distorting the supply-side of the electricity market by increasing the short-run costs of generation. Section 5 describes our back-of-the-envelope calculations linking reductions in generating-unit outages to increases in equilibrium quantities of electricity. Section 6 concludes.

2 Background and Data

In 1948, India's Electricity Supply Act gave rise to State Electricity Boards (SEBs), responsible for regulating electricity generation, transmission, and distribution. The electricity prices set by these SEBs were too low to recover costs, leading to supply shortages, a lack of investment in generation capacity, and bankruptcy among state electricity companies. In response, the Indian electricity sector was charged with reform, beginning with the introduction of State Electricity Regulatory Commissions, starting in 1996.

3. 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.

The most significant reform in recent history has been the 2003 Electricity Act, which called for an overhaul of the Indian power market. The stated goal of the Act was to facilitate competition in the supply of electricity by changing many different aspects of the industry, including opening up access to electricity transmission grids as well as removing the licensing requirement for the generation and distribution of electricity. In the wake of the Act, private investment in generation increased dramatically.

2.1 Capacity and Generation

The Central Electricity Authority (CEA) monitors the operations of all utility-scale fossil, hydroelectric, and nuclear power plants in India.[4] This includes centrally owned, state-owned, and privately owned facilities—each of which must report their daily operational capacity, scheduled generation, and actual generation to the CEA. The CEA then publishes Daily Generation Reports, which we have digitized for each day from January 1, 2013 to March 31, 2017.[5] These reports include data on 485 unique plants that represent 291 GW of India's 330 GW of electric generating capacity, with an average total generation of approximately 2.9 TWh per day.

[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). Figure 2 shows the locations of these power plants across India.

[Figure 2 about here]

4. Wind and solar resources instead fall under the Ministry of Renewable Energy. We know of no publicly available daily generation data for non-hydro renewables, which comprise less than 20 percent of India's total generation capacity. These technologies are non-dispatchable and have extremely low marginal cost. Utility-scale plants have capacity greater than 50 MW.

5. We are in the process of extending this sample through December 31, 2017.

The CEA provides detailed information on plant operations in the annual *Review of Performance of Thermal Power Stations*, including operating heat rates for a large subset of coal-fired plants. Heat rates are the standard metric of electric generators' (inverse) thermal efficiency, measured here in kilocalories per kilowatt-hour. We digitized the 2012–2014 *Reviews* (the most recent available years), and we also obtained the 1997–2009 data from Chan, Cropper, and Malik (2014).[6] Since our analysis spans 2013–(early) 2017, we assign each plant its most recent heat rate that we observe.[7] By combining CEA heat rates with data from the CEA's Monthly coal Reports, we are also able to assign plant-specific coal grades (i.e. coal's gross calorific value in kcal per kilogram). 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]

Most fossil fuel plants procure fuel on long-term contracts called Fuel Supply Agreements (FSAs). Typically, plants must sign an FSA before they are allowed to begin operations. The majority of coal-fired plants purchase coal from Coal India, Ltd., the government-owned coal supply monopoly. Nearly all such plants buy their coal at grade-specific prices set by the Ministry of Coal.[9] We collect coalfield-level prices from the Ministry of Coal, which then assign to individual power plants based on coal grade and geographic proximity.[10] Transportation costs also contribute a substantial share of the total fuel costs paid by coal plants, and the vast majority of Indian coal travels by rail. We approximate plant-specific transportation costs per kilogram of coal using plants' distance from coalfields and

6. We thank the authors for sharing these data.

7. For only 16 plants appearing in the *Reviews*, this 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.

8. Centrally owned and state-owned plants face stricter reporting requirements than private power producers (Chan, Cropper, and Malik (2014)).

9. These are regulated "pithead" prices, and they do not include the cost of transporting the coal from mines to plants. After our sample period, the government introduced the "Scheme to Harness and Allocate Kolya (Coal) Transparently in India" policy (a.k.a. Shakti), which allocates *new* coal contracts to generating units based on an auction mechanism. The first auction ran in September 2017, after the end of our current sample.

10. We calculate the distance between coal plants and coalfields, combining hand-coded plant latitude/longitude with geospatial data on India's coalfields from the U.S. Geological Survey.

coal-specific freight rates from the Freight Operations Information System of the Ministry of Railways.[11]

For natural gas-fired power plants, we assign heat rates based on the CEA's Monthly Gas Reports from 2012, 2016, and 2017.[12] These data enable us to assign 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*, and assign these prices as plants' fuel costs.

Combining heat rates and fuel costs, we are able to approximate plants' marginal cost of generation for fossil plants.[13] We plot the resulting aggregate supply curve in Figure 3, assigning the 7 nuclear plants marginal costs based on assumptions of the Ministry of Power. Importantly, this aggregate supply curve omits hydro generation, whose marginal costs are close to zero but difficult to characterize across time.[14]

[Figure 3 about here]

One salient feature of Indian power plants is their low technical efficiency (Chan, Cropper, and Malik (2014)): a substantial share of total capacity is typically unavailable to generate on any given day. Each day, the CEA reports MW of capacity for each plant that is under outage (i.e. unavailable to generate). These outages are classified as either "planned" or "unplanned": planned outages are scheduled far in advance of production, while unplanned outages reflect (supposedly) unanticipated reductions in capacity.[15] Figure 4 plots the share of capacity in our sample that is under outage, as a daily time series. On average, over a quarter of generation capacity—27 percent—is unavailable each day. The majority of these outages are classified as "unplanned", as "planned" outages represent at most 30 percent of the total unavailable capacity on a given day. We also construct our own classification for

11. We also incorporate royalties and other taxes into plants' all-in coal procurement costs.

12. These are the only available years, and we assign each plant its average observed heat rate. We follow the Ministry of Natural Gas and Petroleum in assuming 10,000 kCal per standard cubic meter of natural gas.

13. We assume marginal costs are constant within each generating unit, for all levels of production up to the unit's strict capacity constraint. This is a standard assumption in the literature on electricity generation.

14. Dispatchable hydro generators face a complex dynamic optimization problem, as generation today may come at the expense of generation tomorrow due to a finite supply of water. Non-dispatchable run-of-river hydro (along with wind and solar) enters the supply curve at (virtually) zero marginal cost.

15. Note that "planned" outages are required to be reported to the grid operator 365 days in advance.

outages that are "internal" to the plant, in order to distinguish between outages potentially within a plant's control (e.g., maintenance, operational issues) vs. outages caused by external factors (e.g., transmission failures, fuel shortages).[16]

[Figure 4 about here]

2.2 Electricity Distribution

Electricity distribution companies, known as "discoms", purchase the majority of electricity sold by Indian power plants. Discoms then resell electricity to consumers at regulated prices, which are set by either state or federal regulatory commissions. However, in an effort to ensure affordable power for residential consumers, these prices are often too low for discoms to recover the costs of purchasing and distributing electricity. Compounding this problem are low payment rates, whereby a substantial share of customers do not pay their bills. As a result, most discoms face serious financial difficulties and require subsidies from state governments. Even inclusive of these subsidies, discoms in many states do not earn positive profits (Pargal and Banerjee (2014); Central Electricity Regulatory Commission (2018b)).

Discoms have responded to these financial difficulties by simply failing to meet electricity demand in all hours and locations. Rolling blackouts (a.k.a. "load shedding") are common. Most Indian states incur annual electricity supply deficits, such that consumers receive less power in the aggregate than forecast supply benchmarks (Central Electricity Authority (2018)). Since discoms lose more money on the margin when electricity generating costs are higher, they should be more likely to withhold power from consumers when wholesale electricity prices are high.

2.3 Long-term Contracts and the Short-term Market

Nearly 90 percent of electricity in India is sold via long-term contracts between electricity producers and distribution companies. Most bilateral contracts (a.k.a. power purchase

16. Unlike deregulated U.S. electricity markets where plant outages often reflect the exercise of market power, the vast majority of Indian power plants receive a fixed contract price per kWh and have no means of increasing this price by withholding production (see Section 2.3 below). Hence, it is extremely unlikely that market power is driving India's high outage rate.

agreements, or PPAs) take the following form. First, the contract enumerates the set of electricity generating units considered. Then, the buyer (i.e. discom) will claim a certain percentage of each unit's capacity to be under contract. This obliges the seller (i.e. generating unit) to allocate this share of its production exclusively to its contracted buyer; however, the buyer is *not* obliged to purchase each unit of power generated under this exclusivity agreement. Contracts also assign each generator a plant load factor (PLF), or a incentivized production target as a share of total potential production. For example, if a unit with 100 MW of capacity is contracted for 100 percent of its capacity with a PLF of 85 percent, then it is obliged to deliver 744.6 GWh each year (i.e. 100 MW $\times$ 8760 hours $\times$ 0.85).

Contracted generators typically receive a price per kWh that is constant across all hours of day and times of year.[17] Hence, they lack any dynamic incentives to operate at times when the value of power is high, or to shift outages to times when the value of power is low.[18] Financial trading is not allowed in Indian power markets, meaning that plants tied up under long-term contracts cannot hedge against market risks, or pay lower-cost plants to generate in their stead. Moreover, discoms have the right to recall until 90 minutes prior to the market realizing, leaving contracted generators effectively no opportunity to sell (eventually uncalled) capacity in the short-term or day-ahead markets, if that capacity is contracted.

Short-term transactions make up the remaining 10 percent of Indian electricity sales. Approximately 5 percent is traded on short-term bilateral contracts with a duration of less than 1 year, while 3–4 percent is traded on one of two short-term power exchanges: the Indian Electricity Exchange (IEX) and Power Exchange India (PXIL).[19] These exchanges run day-ahead uniform-price auctions, where generators submit supply offer curves, discoms submit demand bid curves, and the market clears by aggregating supply and demand. The exchanges clear separately for each 15-minute period, for each of India's five transmission

17. This volumetric price has two components: a "fixed cost" payment per kWh meant to cover plants' fixed costs of operation and long-term financing; and per-kWh payment meant to cover plants' variable costs of generation. Both payment rates are functions of the regulator's assessment of each plant's fixed and variable costs. When a contracted plant stands ready to sell, but the discom exercises its right not to buy, the discom must still pay the "fixed cost" payment per kWh (*as if* the plant had generated).

18. Planned outages are agreed upon a year in advance. Plants cannot use planned outages to avoid operating in high-value periods, and discoms cannot ask the plant to shift their planned maintenance to *ex post* low-value periods.

19. The remaining transactions occur through the "deviation settlement mechanism", which balances deviations from scheduled generation in real-time.

regions: North, Northeast, East, West, and South (see Figure 2). Due to the one-way restrictions of bilateral contracts, as well as the absence of financial trading, generators are only able to submit offers to the exchanges for capacity that is *not* contracted (Central Electricity Regulatory Commission (2018a)). In practice, this precludes most plants from actively participating in the day-ahead market, limiting the size of the exchanges.

We have collected aggregate supply and demand curves for both exchanges. For IEX, this involved digitizing 20,298 15-minute intervals worth of supply and demand curves, spanning 808 days.[20] Figure 5 shows an example of two such intervals, with curves that we have converted to Cartesian coordinates. We have also collected data on posted IEX market-clearing prices and quantities, at a 15-minute resolution for all five regions. Across our sample, the average IEX market clearing price was 2,976 Rs/MWh, while the average volume cleared was 3,934 MWh per 15-minute interval.[21]

[Figure 5 about here]

3 Price Elasticity of Wholesale Electricity Demand

This section quantifies the price elasticity of demand for electricity in the Indian wholesale electricity market. We begin by demonstrating that demand for electricity traded on the IEX, the final market for power in India, is quite elastic. We then use an instrumental variables approach to show that *total* quantity demanded is also responsive to the IEX price. This demonstrates that aggregate demand responds to IEX prices, which guards against the possibility that contracted capacity is able to fully substitute for power sold on the IEX.

3.1 Electricity Demand in the IEX

To quantify the price elasticity of demand in the IEX, we simply use aggregated demand curve data from the IEX. As described in Section 2, we have digitized tens of thousands of individual JPEG images of IEX supply and demand curves.[22] The top left panel of Figure 6

20. We are still in the process of digitizing the remaining intervals, which is computationally intensive. Our current dataset is approximately a 10 percent random sample of intervals.

21. For now, we focus our analysis on the IEX, the larger of the two markets.

22. Digitization of the remaining data is in progress.

displays data on IEX demand curves. The gray lines in this panel show average demand curves for each 15-minute interval of the day, generated by averaging quantity demanded over each interval in 5 Rs-wide price bins. The thick blue line shows mean demand over all intervals. The top right panel presents the average demand elasticity for each interval as a function of distance from the market clearing price.[23] The bottom panel shows a histogram of market clearing prices across all date-interval pairs in our dataset.

[Figure 6 about here]

Figure 6 highlights two important facts about the IEX data. First, demand for electricity in the IEX is highly elastic, in stark contrast with the completely inelastic wholesale electricity demand in the West. The average elasticity at the IEX market clearing price is -1.5 across intervals, quite far from zero.[24] Second, IEX demand curves extend to prices *below* the average market clearing price, meaning that expanding the IEX supply curve could meaningfully increase the equilibrium quantity at the margin.

In order for elastic IEX demand to impact retail electricity consumers, it must be the case that bids from discoms—rather than other IEX participants—make up the elastic portion of the IEX demand curve close to the market clearing price. While other buyers besides discoms purchase power on the IEX (e.g. manufacturing plants), we have three pieces of evidence that discoms' bids constitute the elastic portion of the IEX demand curve. First, Ryan (2017) has access to anonymized bid data which makes up the aggregate IEX demand curve (albeit for an earlier time period), and notes that "the large, flat steps in the demand curve are bids of electricity distribution companies". Second, analogous data from the PXIL (the smaller of India's two power exchanges) also exhibits highly elastic demand, and its buyers are nearly all discoms. Third, officials from the IEX confirmed that this was the case in a recent meeting.[25]

While the IEX makes up only a small fraction of the overall market, shift out the IEX supply curve has the potential to influence the quantity of blackouts (or "load shedding").

23. We construct these mean elasticities by first smoothing the demand function in 100 rupee units, then computing the elasticities for 5 rupee changes in price, and finally computing average elasticities across dates for each interval's price bins.

24. Without prior smoothing, these elasticities can be as large as -8 around the market clearing price.

25. We have requested anonymized bid data from the IEX, in order to confirm this ourselves.

According to data from POSOCO, the Indian grid operator, the average daily peak shortage is 3,561 MW. This is a similar magnitude to the amount of power offered by suppliers in the IEX.[26] In an average interval, there is 707 MW of demand at prices below the market clearing price, or 20% of the peak shortage, suggesting that lowering the IEX price and/or shifting out IEX supply has the potential to meaningfully increase total quantity supplied.

3.2 Instrumental Variables Approach

The IEX represents a small portion of total electricity sales, over which the discoms' demand for power is quite elastic. In this section, we estimate the price elasticity of *aggregate* electricity demand, to see whether discoms also reduce *aggregate* power purchases when the IEX price increases. We regress (log) aggregate electricity quantity on (log) price in the day-ahead IEX market—the last market clearing price before the market realizes. To overcome simultaneity bias, we need to instrument for price using a supply shifter. One such shifter is the fraction of capacity that is on unplanned outage, which should impact electricity supply but not electricity demand.

For each Indian electricity region r, we estimate the following system of time series regressions using two-stage least squares:

$$\log\left(Q_{mdt}\right) = \beta \log\left(\hat{P}_{mdt}\right) + \gamma \mathbf{X}_{mdt} + \eta_m + \delta_d + \varepsilon_{mdt} \tag{1}$$

$$\log\left(P_{mdt}\right) = \nu[\text{ShareUnplannedOutage}]_{mdt} + \gamma \mathbf{X}_{mdt} + \eta_m + \delta_d + \omega_{mdt} \tag{2}$$

where Q_{mdt} is the aggregate quantity of electricity generated on date t, in month m, on day-of-week d; and P_{mdt} is the IEX market clearing price, average across all 96 intervals on each date t. Using aggregate quantity on the left-hand side (rather than IEX quantity) guards against the possibility of discoms substituting IEX purchases with contract purchases, or vice versa. It also allows us to estimate how IEX price affects the full market, beyond the 3–5 percent of demand met by IEX supply. $[\text{ShareUnplannedOutage}]_{mdt}$ is the fraction

26. This "shortage" measure is defined as the difference between the quantity of power that demanding entities ask POSOCO to provide to them vs. what they actually receive. Note that this is net of additional load shedding: if a discom faces total customer demand of 100 MW, but wishes to provide only 75 MW of power, the shortage number does not reflect this 25 MW gap.

of generation capacity that is under a *unplanned* outage. In constructing this instrument, we remove capacity from both the numerator and denominator with marginal costs higher than the market clearing price; this improves the strength of the instrument by not counting outages in high-cost peaker plants, which rarely generate.[27] The matrix $\mathbf{X}_{mdt}$ controls for state-specific time series of daily maximum temperatures, as temperatures could impact both demand and unplanned outages. We also include month-of-sample fixed effects (η_m) and day-of-week fixed effects (δ_d). We estimate Newey-West standard errors with a 7 day lag.

Table 1 reports results from estimating Equations (1)–(2) separately for each region. For the North, Northeast, East, and West regions, we find a strong first stage with positive and statistically significant point estimates. In these four regions, when low-cost plants have more unplanned outages, the IEX market clearing price increases as expected.[28] However, the first stage has the opposite sign for the South region, with an extremely low F statistic of 1.17. This likely reflects the spatial configuration of the grid: while the North, Northeast, East, and West regions frequently import power, the South region almost always exports. Hence, the unit setting the South regions' IEX market clearing price is unlikely to be located in the South, thereby breaking the correlation between own-region outages and own-region prices. For the other four regions, we find strong positive relationships between unplanned outages and IEX price, with F statistics ranging from 39.43 to 264.11.

[Table 1 about here]

Next, we estimate aggregate demand elasticities via two-stage least squared. For the four regions with strong first stages, we can reject perfectly inelastic aggregate demand at the the 1 percent level. Our estimated elasticities are -0.19 in the North region, -0.15 in the Northeast region, -0.20 in the East region, and -0.14 in the West region. These

27. Our results are similar if we omit plants with marginal costs above: (i) mean IEX price on date t; (ii) maximum IEX price on date t; (iii) mean (or maximum) IEX price on date t plus one standard deviation in IEX price; (iv) lagged mean IEX price on date $t-1$; or (v) mean IEX price over days $t-7$ to $t-1$. Moving from (i) to (v) reduces the instrument's predictive power, while also decreasing the likelihood that this selection criterion violates the exclusion restriction. Point estimates are relatively stable across (i)–(v), consistent with the instrument shifting only supply. Table 1 presents results using selection criterion (i).

28. We are currently working to strengthen this first stage even further, by (i) isolating outages among IEX participating plants, and (ii) instrumenting with observed shifts in IEX supply curves.

estimates may even be lower bounds on discoms' true elasticities, since they do not account for greater responsiveness during peak (as opposed to off-peak) periods. They also imply that elastic IEX sales do not simply crowd out long-term-contract purchases one-for-one. Rather, at higher IEX prices, discoms purchase less *aggregate* electricity—meaning that less power reaches consumers.

4 Plant Outages

Section 3 documents that wholesale electricity demand is downward sloping in India. It follows that any supply-side distortions that increase the equilibrium price in the wholesale market *also* reduce the equilibrium quantity of electricity that discoms buy and resell to consumers. In this section, we investigate the extent to which Indian wholesale electricity supply exhibits short-run allocative inefficiencies. We compare the total variable costs of generation as observed in the market vs. a counterfactual "least-cost" scenario where plants are dispatched in order of lowest-to-highest cost. The difference in total variable costs is quantitatively large, but greatly attenuates once we incorporate plant outages into the least-cost counterfactual. Then, we provide descriptive evidence that both the frequency and timing of plant outages are not explained by economic factors such as plant age, marginal costs, or temperature. Taken together, this suggests that plant outages represent a significant supply-side distortion that effectively reduces the aggregate quantity of electricity sold in the wholesale market.

4.1 Costs under Observed vs. Least-Cost Dispatch

We calculate the aggregate variable costs of electricity production for each day t in our sample. This calculation simply multiplies unit i's observed generation Q_{it}^{OBS} by its marginal cost MC_{it}. Summing across all units gives us the total observed variable costs:

$$TC_t^{OBS} \equiv \sum_i MC_{it} Q_{it}^{OBS} \tag{3}$$

We assume constant marginal costs within each unit over feasible levels of generation, from $Q_{it} = 0$ up to unit i's strict capacity constraint $Q_{it} = \overline{Q}_i$; this assumption is standard in the literature on electricity supply.[29]

Next, we repeat the same daily calculations under a "least-cost" counterfactual. This approach re-dispatches all generating units in order of lowest-to-highest cost, until aggregate counterfactual generation equals aggregate observed generation on day t.[30] This is equivalent to the following cost-minimization problem:

$$TC_t^{LC} = \min_{\{Q_{it}^{LC}\}} \sum_i MC_{it} Q_{it}^{LC} \quad \text{s.t.} \quad Q_{it}^{LC} \in [0, \overline{Q}_i] \;, \quad \sum_i Q_{it}^{LC} = \sum_i Q_{it}^{OBS} \;\; \forall \, t \tag{4}$$

We do *not* interpret this re-dispatching of plants as a feasible best-case scenario, as it dramatically simplifies the technical constraints associated with electricity generation and transmission. Rather, this idealized benchmark helps us characterize the relative importance of factors that might contribute to high marginal costs of generation.

The first row of Table 2 presents the daily average "cost difference", which we calculate as:

$$TC_t^{DIFF} \equiv TC_t^{OBS} - TC_t^{LC} = \sum_i MC_{it}(Q_{it}^{OBS} - Q_{it}^{LC}) \tag{5}$$

Averaging across 1,530 sample days, the mean cost difference is 4.3 million U.S. dollars per day. The minimum and maximum daily cost differences are \$2.8 million and \$5.9 million respectively (reported in brackets), which reflects substantial variation in both least-cost and observed costs. The right-most column reports the cost differences in percentage terms (i.e., TC_t^{DIFF}/TC_t^{OBS}), and finds an average daily cost difference of 16 percent.

Transmission constraints offer one potential explanation for this 16 percent cost difference between observed vs. least-cost dispatch. Electricity can only flow between regions of the grid if unconstrained transmission capacity is available, yet Equation (4) ignores transmission when constructing counterfactual least-cost dispatch. To account for constraints between India's five transmission regions (see Figure 2), we recompute Equation (4) separately

29. Section 2 describes how we construct marginal costs by combining data on fuel costs and heat rates.

30. We assign each unit's capacity $\overline{Q}_i$ as the 98th percentile of its observed generation during our sample period. Our results are not sensitive to this assumption, and the ensuing calculations are quite similar using the 80th percentile of observed generation.

for each region and sum total costs across all five regions. This eliminates counterfactual re-dispatching of plants *across* regions, which may be unfeasible in practice. We report these results in the second row of Table 2, where the average daily cost difference falls from 16.0 to 13.9 percent of observed costs. This implies that most of the cost difference reflects potential re-dispatching *within* transmission regions, rather than unrealistic *between*-region re-dispatching.[31]

Thus far, our cost difference calculations have ignored within-day variation in electricity production. However, the Indian electricity market produces substantially more power during afternoon peak hours, compared to off-peak hours. While we do not directly observe individual plants' hourly generation, we do observe each region's total daily generation separately for peak vs. off-peak periods. This allows us to split observed generation at the plant-by-day level into (approximate) peak/off-peak shares, weighted by the share of daily generation coming in peak/off-peak periods. It also allows us to re-dispatch plants at the sub-daily level, separately for peak vs. off-peak periods.[32] The third row of Table 2 reports results under regional autarky *and* separate peak vs. off-peak re-dispatching. This yields cost differences that are almost identical to the regional autarky scenario with daily re-dispatching. While our weight-averaging approach cannot capture the full within-day variation in generation, these results provide suggestive evidence that within-day variation is not a major driver of the 14 percent cost difference.

Finally, we incorporate plant outages into our counterfactual least-cost dispatch. Section 2.1 discusses how a substantial share of Indian generating capacity is unavailable on most days (see Figure 4). However, our least-cost calculations have thus far assumed that *all* generating capacity is available for re-dispatching on each day. The fourth row of Table 2 builds on the above scenario by taking plants' declared outages as given—that is, least-cost counterfactuals may only re-dispatch capacity that is available (i.e. *not* on outage) on day t. Removing capacity under outage decreases the average daily cost gap from \$3.7 billion to

31. These results are similar if we impost sub-regional autarky, which limits re-dispatching to be within each of India's 13 sub-regions.

32. For example, suppose total generation in the Western region is split 70/30 between peak/off-peak periods on a given day. Then, for each plant in the West region on that day, we assign 70 percent of its observed generation as "peak" and 30 percent of its observed generation as "off-peak". For least-cost dispatch, we allow plants 6 hours of potential peak generation and 18 hours of potential off-peak generation.

$1.2 billion, or from 13.6 percent to 5.1 percent of average observed costs. This implies that the above cost differences are predominantly explained by unavailable low-cost generating capacity. In fact, the minimum cost difference in this scenario is zero: on a few sample days, removing plant outages from least-cost dispatch yields the same total variable costs as observed dispatch. Figure 7 plots kernel densities of daily cost differences for all four of the above scenarios, underscoring the extent to which accounting for plant outages shifts the distribution towards zero.

[Table 2 and Figure 7 about here]

Again, we emphasize that these calculations do not fully characterize short-run inefficiencies in Indian electricity provision. A variety of unmodeled (or imperfectly modeled) factors might yield differences between least-cost and observed costs in the absence of short-run market distortions.[33] Even so, the importance of plant outages is striking—the magnitude of the cost difference hinges on whether we take as given plants' available capacity on each day. On the one hand, large power plants require periodic maintenance and may face unanticipated operational disruptions; hence, an accurate least-cost counterfactual should not assume all generating units are always available to produce at capacity.[34] On the other hand, India's high outage rate suggests the presence of excessive outages that do *not* reflect plants' technical constraints; eliminating any such excessive outages could yield substantial decreases in the total variable costs of Indian electricity supply.

4.2 Frequency and Timing of Plant Outages

Having established that plant outages explain a substantial share of the gap between observed vs. least-cost dispatch, we now investigate the frequency and timing of these outages. It is obviously unrealistic to expect utility-scale power plants to make their full capacity available

33. Such unmodeled factors might include localized transmission constraints around load pockets, or plants' dynamic operating constraints (i.e. startup and ramping).

34. Recall that 90 percent of generating capacity is committed to long-term contracts, where buyers pay a constant rate per kWh. Receiving a fixed price per kWh eliminates any incentive to withhold capacity, implying that India's high outage rate is *not* indicative of the wide-spread exercise of market power. In fact, the typical long-term contract provides weak incentives for plants to avoid outages, in the form of small "fixed cost" payments contingent on capacity being *available* to generate.

to generate 100 percent of the time. Thermal generators have technically complex operations, require periodic maintenance, and face numerous engineering constraints that might preclude generating at full capacity. However, given the high aggregate outage rates shown in Figure 4, many Indian plants should be technically capable of reducing the amount of MW-days under outage.

Figure 8 summarizes the frequency of outages across all 170 coal plants in our sample. The top panel reveals substantial dispersion across plants: a quarter of coal plants have total outage rates greater than 44 percent, meaning that over 44 percent of these plants' capacity is unavailable on an average sample day. By contrast, the "best-performing" quartile of coal plants have total outage rates less than 15 percent. The bottom panel summarizes the frequency of "internal" outages, or outages attributed to factors within the physical footprint of each power plant (and thereby more likely to be within the plant's control than outages caused by external factors such as transmission failures). We see the "best-performing" quartile of plants have internal outage rates less than 9 percent, while the "worst-performing" quartile of plants have internal outage rates over 25 percent. Importantly, this variation in outage rates is *not* explained by plant-specific covariates such as vintage, total capacity, region, heat rate, marginal cost, or ownership sector (i.e. federal, state, or private).

Aside from the *frequency* of outages, most Indian power plants lack the incentive to shift discretionary outages from high-value periods to low-value periods. This is because plants under long-term contracts receive a price per kWh that is time-invariant. These plants tend to be effectively locked-out of day-ahead markets, where the market-clearing price would provides a time-varying incentive of *when* to most profitably generate. Table 3 tests whether plants are less likely to have capacity under outage on days with high temperatures, a proxy for days with a high marginal value of electricity. Each column estimates the following regression:

$$[\text{ShareOutage}]_{ist} = \beta\,[\text{MaxTemperature}]_{st} + \gamma_i + \delta_t + \varepsilon_{it} \tag{6}$$

The dependent variable is the share of plant i's capacity on outage of a given type on day t, while $[\text{MaxTemperature}]_{st}$ is the spatial average maximum temperature in state s on day

t. We control for plant fixed effects (γ_i) and month-of-sample fixed effects (δ_t), and cluster standard errors at the district level (with 134 clusters).

[Table 3 about here]

We find only very weak associations between outages and temperature. Even though point estimates are negative and statistically significant for total and unplanned outages (Columns (1) and (3)), the magnitudes are quite small—a 10°C increase in maximum temperature only translates to a 4 percentage-point decrease on plants' share of capacity under outage.[35] Most strikingly, this small correlation disappears in Column (5), which only considers outages internal to each plant. This suggests that better dynamic incentives could yield more optimal timing of outages, as plants' should be more willing to generate on hot days if they expect to receive a higher price per kWh.

We do not (currently) observe which plants participate in the IEX versus which plants are contracted for 100 percent of their capacity.[36] While theory suggests that plants which participate in the IEX face stronger incentives to supply more capacity on high-demand days than their non-IEX counterparts (since IEX plants will capture inframarginal rents which increase in the market clearing prices), in the absence of these data, we cannot test this directly. However, we can perform an indirect test of this by examining the extent to which the aggregate IEX supply curve responds to temperature, the same demand shifter we used to look at capacity under outage in the full market.

To do this, we regress:

$$\log(\text{Offered volume})_{rt} = \beta\,[\text{MaxTemperature}]_{rt} + \gamma_r + \delta_t + \varepsilon_{rt} \tag{7}$$

Where $\log(\text{Offered volume})_{rt}$ is the log of the quantity of electricity bid into the IEX by suppliers in region r during day t. We estimate five specifications, ranging from no fixed effects to a specification with region-by-year and month-of-sample fixed effects. We cluster our standard errors at the month-of-sample level (with only five regions, clustering at the

35. Results are similar if we specify a 2nd- or 3rd-order polynomial in maximum temperature.
36. We are in the process of negotiating an agreement with the IEX in order to acquire these data.

region level would lead to downward biased standard errors). Table 4 presents the results of these regressions.

[Table 4 about here]

In contrast to the results from outages in the full market presented in Table 3 above, we find relatively strong links between temperature and IEX supply. Across specifications, we find a statistically significant relationship between temperature and IEX supply. These effects are economically meaningful: a 10 C increase in temperature are associated with between 11 and 61 increases in quantity offered by IEX plants. This suggests that IEX plants are meaningfully responding to a demand shifter, unlike the market as a whole (with the important caveat that quantity offered is not the same object as the share of capacity under outage).

5 Back-of-the-envelope calculations

In this section, we discuss our preliminary back-of-the-envelope approach to calculating the increase in quantity supplied that could arise from improving power plant incentives and reducing outages, as a result of elastic demand for electricity in the wholesale market. We proceed in several steps, intended to capture the effects of marginal changes in power plant availability on total quantity produced.

First, for each date in our sample, we identify the amount of capacity that is on outage at plants whose marginal cost is below $p \times \text{MCP}$, where MCP is the market clearing price in the IEX and $p \in [0, 1]$. These are the outages that are marginal to the IEX; outages at plants with prices above the market clearing price would not affect total quantity demanded. We implement this approach with a conservative value of $p = 0.5$, rather than simply setting $p = 1$, to account for the possibilities that (i) our marginal cost estimates are likely to be somewhat downward biased, since we do not observe non-fuel operating costs, and (ii) plants in the IEX face an incentive to effectively charge a markup (typically done by withholding quantity). Ryan (2017) estimates this effective markup to be 20%.

Next, we add $r \in [0, 1]$ fraction of these outages into the IEX supply curve, mimicking an increase in quantity supplied that would result from increased incentives to be available and capture inframarginal rents in the wholesale electricity market. In these preliminary estimates, we only consider overall reductions in outages. This is consistent with the narrative of Davis and Wolfram (2012), in which deregulation and consolidation of nuclear power plants increased operating performance by 10%, largely through reductions in maintenance. In ongoing work, we are also considering the consequences of keeping the *level* of outages the same, while adjusting the *timing* of these outages from lower-demand to higher-demand days.

Finally, we interact our new, expanded IEX supply curve with the existing IEX demand curve, and calculate the new equilibrium price and quantity. This allows us to quantify the changes in market clearing prices and quantity demanded in the IEX that would result from a reduction in capacity under outage.

Figure 9 illustrates this approach for one example day. The light gray line shows the original IEX supply curve. The blue lines show our expanded supply curves, with the lightest line adding in 10 percent of marginal outage capacity, and the darkest line adding in 100 percent of marginal outage capacity. We define marginal capacity as capacity under outage at plants with marginal cost less than or equal to 50 percent of the IEX market clearing price.

[Figure 9 about here]

Figure 9 demonstrates that expanding IEX supply by returning capacity under outage to the supply curve has the potential to lower market clearing prices and, as a result of downward-sloping demand, increase equilibrium quantities in the IEX. We next implement this approach for all of the days in our sample. Figure 10 plots the results. The left panel of Figure 10 shows the distribution of market clearing prices under the existing supply and demand curves (gray line) and ten counterfactual supply curves, ranging from adding 10 to 100 percent (lightest blue to darkest blue) of outage capacity into the IEX supply curve. The right panel of Figure 10 shows the distribution of market clearing quantities that result from this same exercise. For these results, we assume that plants with marginal costs below 50 percent of the market clearing price are marginal; we do not adjust outages of plants with higher marginal costs.

[Figure 10 about here]

We find that adding capacity that was under outage into the IEX supply curve has a substantial impact on the market clearing prices in the IEX. We find that median market clearing prices are reduced from 1325 Rs/MWh to 1101 Rs/MWh with the addition of 20 percent of low-cost outage capacity. This in turn translates into an increase in equilibrium quantity of nearly 500 MW, or 14 percent of the peak shortage. Given that this is a relatively marginal change in the supply curve, these appear to be somewhat sizeable effects.

There are two important caveats to these results. First, these estimates use the IEX, rather than full market, demand curves. In order to extrapolate these results to effects on load shedding, we will need to extend these calculations to also use our full market demand curves. These additional calculations are in progress. Second, we are currently adjusting the level of outages. We present a range of results in which we vary the fraction of capacity under outage that is returned to the IEX supply curve, but these sensitivities do not reflect effects that result from changing the *timing* of outages. In order to shed light on the impacts of improving temporal incentives for outages, in ongoing work, we are implementing back-of-the-envelope calculations in which we move outages from low-demand to high-demand days, rather than reducing the total quantity of outages.

6 Conclusion

Economic development is strongly linked to increased energy consumption. Frequent blackouts may therefore hamper economic growth and impose undue burdens on poor consumers. In this paper, we demonstrate that Indian utilities have downward-sloping wholesale electricity demand—meaning that as prices rise, they purchase less power and impose rolling blackouts on electricity consumers. Because Indian utilities are price-elastic in the wholesale market, anything that raises marginal generation costs will directly reduce the quantity of electricity sold to retail consumers. We document evidence of elevated variable costs in Indian electricity supply, by comparing observed dispatches with an idealized "least-cost" counterfactual. While we find a cost difference of 13–16 percent, this difference shrinks to 5 percent when we take plants' declared outages as given. Hence, excessive outages appear to be inflating the marginal costs of generation, and we provide evidence that these outages do not covary with economic factors affecting the marginal costs or marginal value of power. Our simple back-of-the-envelope calculations suggest that reducing the outage rate has the potential to reduce peak shortages substantially.

In ongoing work, we are: (1) generating a more complete dataset of wholesale market demand and supply curves, which will allow us to calculate the elasticity of demand for power in the wholesale market; (2) incorporating data from the IEX that identifies which power plants are IEX participants, which will help us refine our instrumental variables estimation for total demand *and* facilitate comparisons of the outage rates of IEX vs. non-IEX plants; and (3) performing more robust counterfactual exercises where we shift out IEX supply or reduce the frequency of plant outages, in order quantify the increase in electricity quantities that could be achieved with reasonable improvements to plants' marginal incentives. Given that roughly 90 percent of electricity generation occurs on restrictive long-term contracts, allowing power plants to more flexibly arbitrage the day-ahead spot market would likely alleviate a meaningful share of the allocative inefficiency in Indian electricity provision.

References

Abeberese, Ama Baafra. 2017. "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.

Borenstein, Severin, James Bushnell, and Frank A Wolak. 2002. "Measuring unilateral market power in wholesale electricity markets: the California market, 1998-2000." *American Economic Review:* 1376–1405.

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.

Central Electricity Authority. 2018. *Load Generation Balance Report, 2018–19.* Technical report.

Central Electricity Regulatory Commission. 2018a. *Discussion Paper on Market Based Economic Dispatch of Electricity: Re-designing of Day-ahead Market (DAM) in India.* Technical report.

———. 2018b. *Report on Short-term Power Market in India: 2017–2018.* Technical report.

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.

Davis, Lucas W, and Catherine Wolfram. 2012. "Deregulation, consolidation, and efficiency: Evidence from US nuclear power." *American Economic Journal: Applied Economics:* 194–225.

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

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.

Hortacsu, Ali, and Steven L Puller. 2008. "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market." *The RAND Journal of Economics* 39 (1): 86–114.

Ito, Koichiro, and Mar Reguant. 2016. "Sequential markets, market power, and arbitrage." *American Economic Review* 106 (7): 1921–57.

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.

Mercadal, Ignacia. n.d. "Dynamic competition and arbitrage in electricity markets: The role of financial players."

Pargal, Sheoli, and Sudeshna Ghosh Banerjee. 2014. *India Power Sector Review: More Power to India: The Challenge of Distribution.* Technical report. The World Bank Group.

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.

Wolak, Frank A. 2003. "Measuring Market Inefficiencies in California's Wholesale Electricity Industry." *American Economic Review:* 425–430.

———. 2007. "Quantifying the supply-side benefits from forward contracting in wholesale electricity markets." *Journal of Applied Econometrics* 22 (7): 1179–1209.

Wolfram, Catherine D. 1999. "Measuring duopoly power in the British electricity spot market." *American Economic Review* 89 (4): 805–826.

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: Aggregate Demand Response to IEX Price

	Aggregate $\log(Q_{\text{region}})$				
	North	Northeast	East	West	South
	(1)	(2)	(3)	(4)	(5)
IV: $\log(P^{\text{IEX}}_{\text{region}})$	-0.190^{***}	-0.150^{***}	-0.203^{***}	-0.140^{***}	2.406
	(0.069)	(0.058)	(0.057)	(0.050)	(2.207)
First stage:					
Share capacity on unplanned outage	0.914^{***}	2.978^{***}	0.641^{***}	1.193^{***}	-0.449^{***}
	(0.115)	(0.160)	(0.075)	(0.135)	(0.116)
Daily temperature controls	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes
Day-of-week FEs	Yes	Yes	Yes	Yes	Yes
Daily observations	1330	1327	1331	1333	406
First-stage F-stat (IV)	58.29	264.11	39.43	72.50	1.17

Notes: Each regression estimates Equations (1)–(2) via two-staged least squares, as a region-specific daily time series. For each regression, the dependent variable is log electricity generation, aggregated by region-day across all reporting plants. Endogenous variables are log-transformed daily IEX clearing prices for each region, averaged across all 96 intervals within each day. The instruments are daily supply shifters: the proportion of each region's generating capacity reporting unplanned outages. Each regression controls for the time series of average daily maximum temperatures for each state constituent to the respective region. Standard errors are Newey-West with 7 lags. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 2: Variable Costs of Electricity Supply

Scenario	Least-cost (M USD / day)	Observed (M USD / day)	Cost Difference (M USD / day)	$100 \times \frac{\text{Difference}}{\text{Observed}}$
National Dispatch	22.8	27.1	4.3	16.0
	[16.5, 28.0]	[20.2, 32.3]	[2.8, 5.9]	[10.1, 22.8]
Regional Autarky	23.4	27.1	3.8	13.9
	[17.1, 28.6]	[20.2, 32.3]	[2.4, 5.1]	[8.5, 19.7]
Regional Autarky + Peak	23.4	27.1	3.7	13.6
	[17.1, 28.7]	[20.2, 32.3]	[2.3, 5.1]	[8.0, 19.6]
Regional Autarky + Peak + Outages	25.9	27.1	1.2	5.1
	[17.6, 32.3]	[20.2, 32.3]	[0, 4.3]	[0, 15.6]
Number of days	1,530	1,530	1,530	1,530

Notes: This table compares the total observed variable costs of electricity generation to total counterfactual variable costs under "least-cost" dispatch. The first column reports the total variable costs under least-cost dispatch (per Equation (4)), while the second column reports total observed variable costs (per Equation (3)), both in millions of US dollars per day. The third column reports the difference between columns 2 and 1 (per Equation (5)). The fourth column divides column 3 by column 2. All columns report averages across 1,530 sample days, with daily minimum and maximum values in brackets. The "National Dispatch" scenario constructs least-cost dispatch according to Equation (4), ignoring plant regions, within-day variation, and capacity under outage. "Regional Autarky" scenarios restrict least-cost re-dispatching to *within* (not *across*) each of India's five transmission regions. "Peak" scenarios re-dispatch generation separately for peak vs. off-peak periods. The "Outages" scenario removes capacity under outage from least-cost dispatch; in other words, it takes plants' declared capacity available as given on each day. See text for further detail. Figure 7 plots kernel densities of the distributions in column 3.

Table 3: Plant Outages and Daily Temperature

	Share of capacity on outage, by category				
	Total	Planned	Forced	Other	Internal
	(1)	(2)	(3)	(4)	(5)
Maximum temperature (°C)	-0.0040^{***}	-0.0002	-0.0040^{***}	0.0002	-0.0002
	(0.0011)	(0.0003)	(0.0009)	(0.0002)	(0.0006)
Plant FEs	Yes	Yes	Yes	Yes	Yes
Month-of-sample FEs	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable	0.328	0.033	0.251	0.044	0.187
Plants	244	244	244	244	244
Plant-day observations	292,653	292,653	292,653	292,653	292,653
R^2	0.480	0.269	0.497	0.725	0.379

Notes: Each regression estimates Equation (6) on a plant-day panel, for a different outage category. We construct dependent variables as the fraction of plant i's total capacity on outage on a given day. Column (1) includes all outages, regardless of the declared justification. Column (2) uses outages classified by the CEA as "planned", which are predominantly planned maintenance periods. Column (3) uses outages classified by the CEA as "forced", which comprise a wide range of declared justifications. The last CEA category is "other", which are largely unplanned and reported in Column (4). Finally, we construct our own outage classification from the CEA's detailed daily outage reports, "internal" outages, which includes all outages internal to the plant (as opposed to outages caused by external factors such as transmission, fuel supply, or discoms). Daily maximum temperatures are spatially averaged at the state level. Regressions are weighted by plants' nameplate capacity. Standard errors are clustered by district, with 134 clusters. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: IEX Volumes and Daily Temperature

	Log(IEX quantity offered)				
	(1)	(2)	(3)	(4)	(5)
Maximum temperature (°C)	0.0607^{***}	0.0113^{**}	0.0106^{**}	0.0249^{***}	0.0274^{***}
	(0.0058)	(0.0049)	(0.0043)	(0.0070)	(0.0065)
Region FEs	No	Yes	Yes	Yes	Yes
Year FEs	No	No	Yes	Yes	Yes
Month-of-sample FEs	No	No	No	Yes	Yes
Region-year FEs	No	No	No	No	Yes
Region-day observations	6,723	6,723	6,723	6,723	6,723
R^2	0.086	0.590	0.625	0.644	0.728

Notes: Each regression estimates Equation (7) on a region-day panel, for a different set of fixed effects. The dependent variable is the log of IEX offered volume. Daily maximum temperatures are spatially averaged at the region level. Standard errors are clustered by month of sample, with 48 clusters. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure 1: Electricity Generation in India – CEA Daily Data

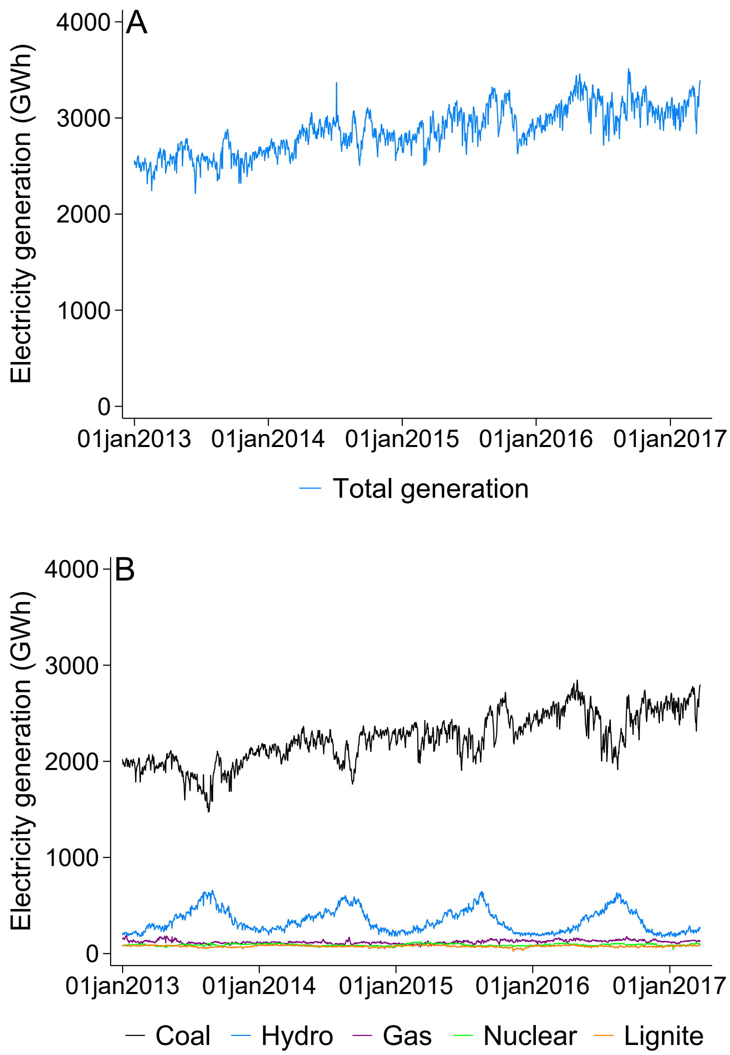

Notes: This figure displays data from the Central Electricity Authority's Daily Generation Reports. Panel A shows total electricity generation in GWh from January 1, 2013 through March 23, 2017. The 485 plants in these data average 2.9 TWh per day. Panel B breaks these plants out by fuel type. The 196 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: Electricity Market Regions of India

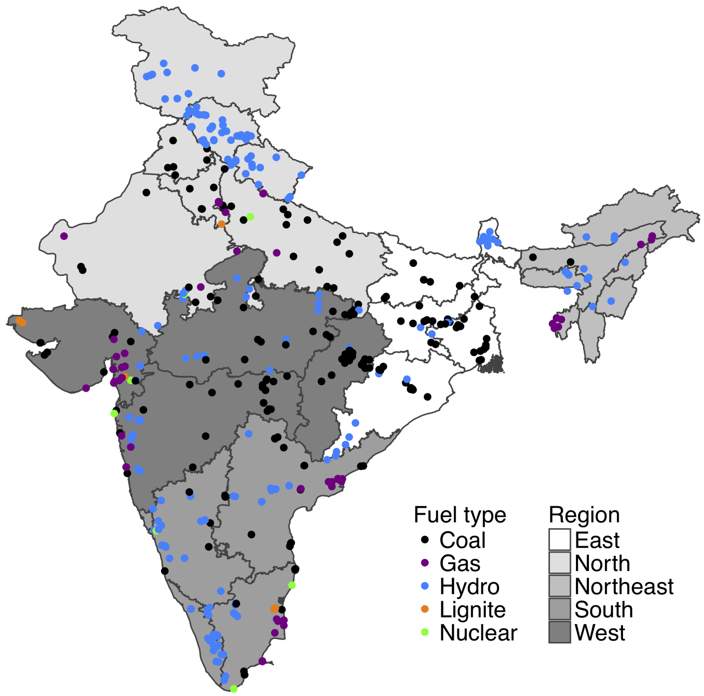

Notes: This figure displays the locations of the 485 power plants that appear in the CEA's Daily Generation Reports. The plant colors indicate the fuel type – coal, gas, hydroelectric, lignite, and nuclear. States are shaded by their major transmission region – North, North East, East, West, and South.

Figure 3: Marginal-cost-based merit order

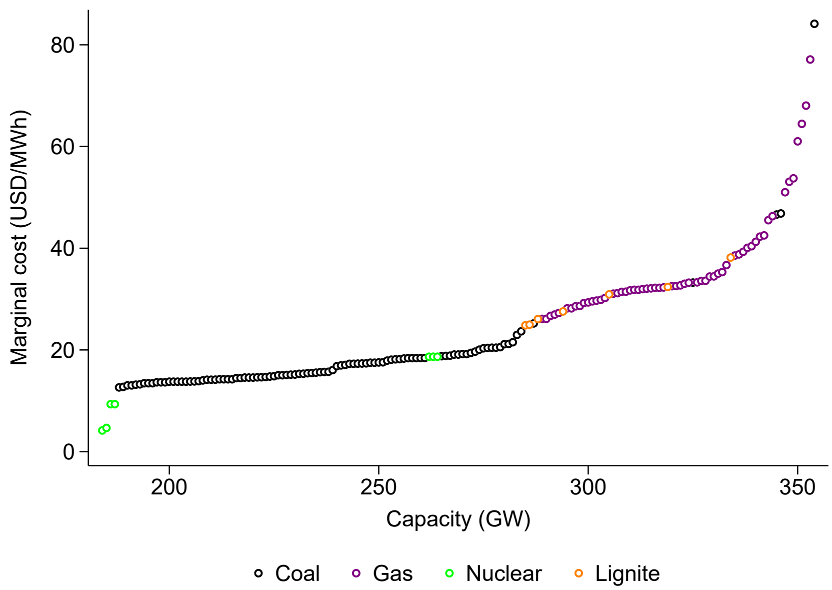

Notes: This figure shows the Indian electricity supply merit order, where we rank plants by their marginal cost according to our own constructed marginal cost estimates. This figure shows all fuel types together, with colors denoting fuel types. Each dot represents 1 GW of capacity; larger plants will 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 4: Time Series of Daily Outages

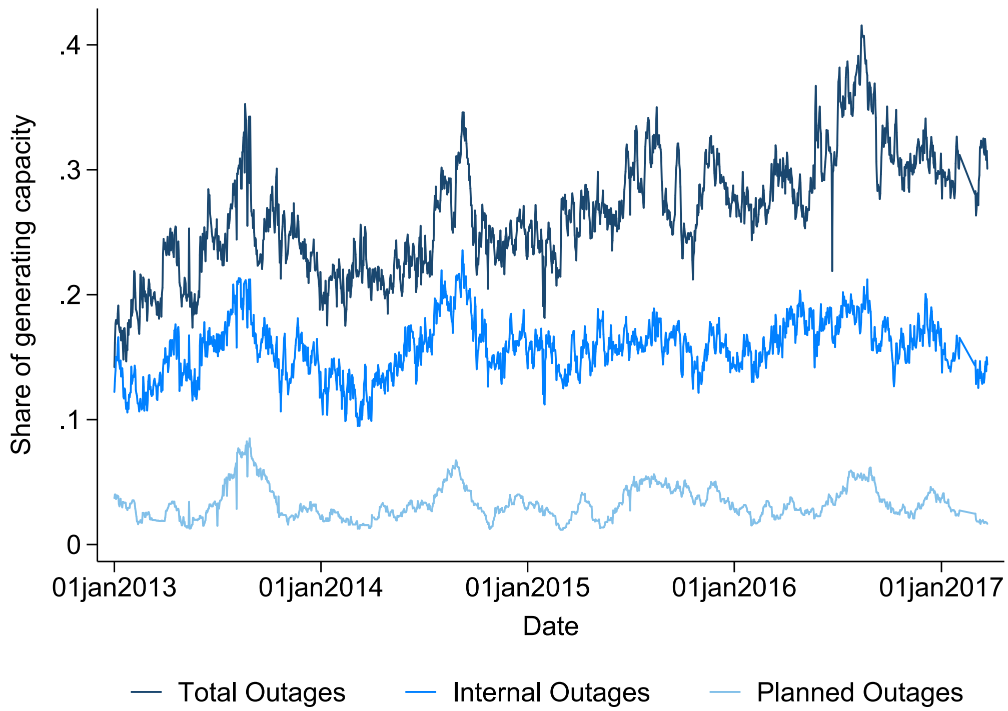

Notes: This figure reports the share of total reporting capacity that on outage (i.e. unavailable to generate) on each each day in our sample. This includes coal, gas, lignite, nuclear and liquid fuels; hydro and renewables capacity is not represented in the numerator or the denominator. The top line divides capacity under outage (for any reason) by total capacity. The middle line includes only outages listed as "internal" to the plant, by our own classification (e.g., maintenance, but not transmission failures). The bottom line includes only outages classified by the CEA as planned.

Figure 5: Example IEX demand and supply curves

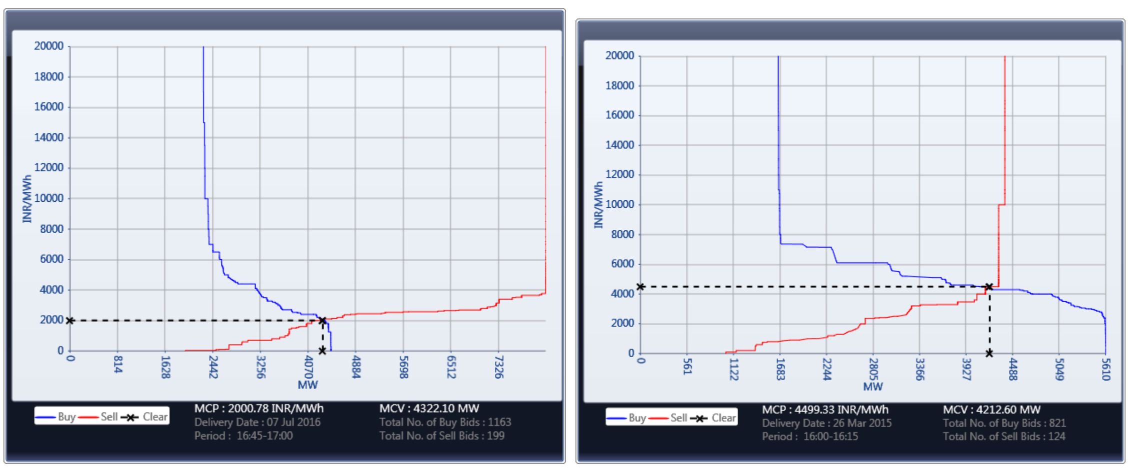

Notes: This figure displays two examples of the raw data we obtained from the Indian Energy Exchange. The image on the left shows the aggregate demand and supply curve for the 16:00-16:15 interval on March 26, 2015. The image on the right shows the same curves for the 16:45–17:00 interval on July 7, 2016. We digitized these images, which originally come in JPEG format, using OCR software.

Figure 6: Elasticity of IEX demand

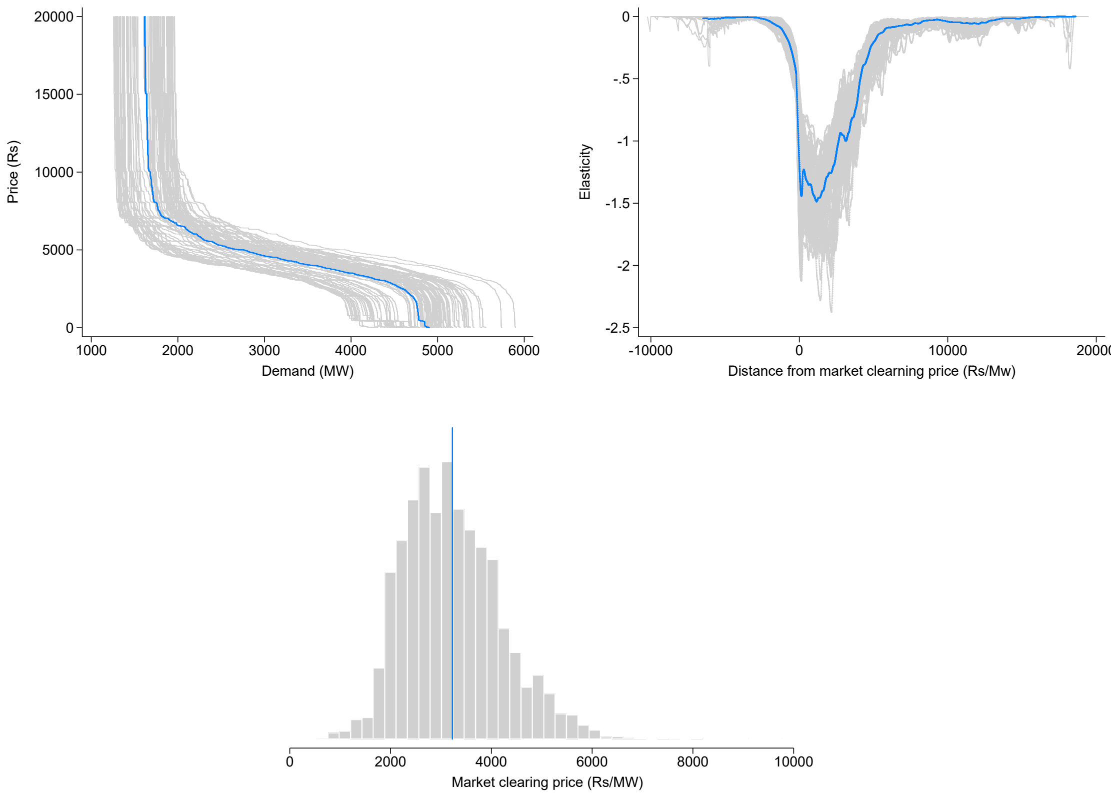

Notes: This figure displays data from the Indian Energy Exchange. The top left panel shows average demand curves for each of the 96 15' intervals of the day. The top right panel shows the price elasticity of demand in the IEX as a function of distance from market clearing price for each of the 96 15' intervals of the day. These elasticities are constructed by smoothing each interval's demand curve over 100 Rs prices, and then computing the elasticity. We then compute and plot the mean elasticity over intervals and distances from the market clearing price. The elasticity is close to -1.5 in the neighborhood the market clearing price. Finally, the bottom central panel shows a histogram of market clearing prices; the mean is just over 3,000 Rs. The right panel Overall, demand for electricity in the IEX is substantially more elastic than the typical perfectly inelastic demand seen in Western power markets.

Figure 7: Distribution of Daily Cost Difference

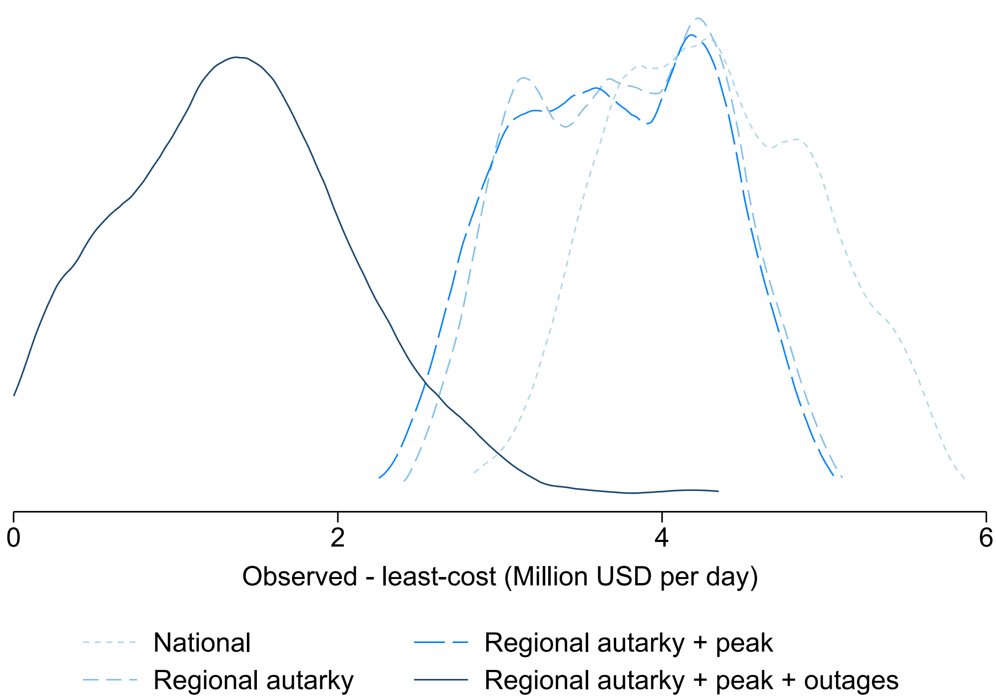

Notes: This figure plots kernel densities of of daily cost differences, for each of the four scenarios reported in Table 2. Column 3 of Table 2 reports the means, minimums, and maximums of these distributions. See the text and notes under Table 2 for further detail.

Figure 8: Histograms of Coal Plant Outage Rates

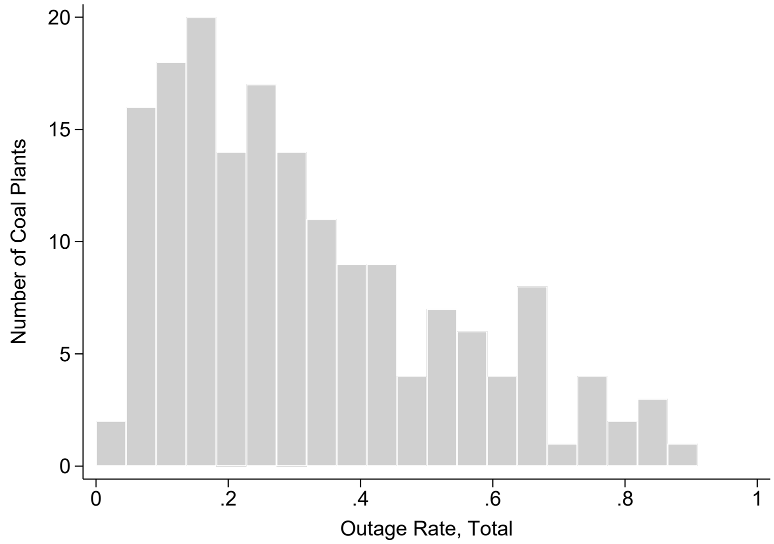

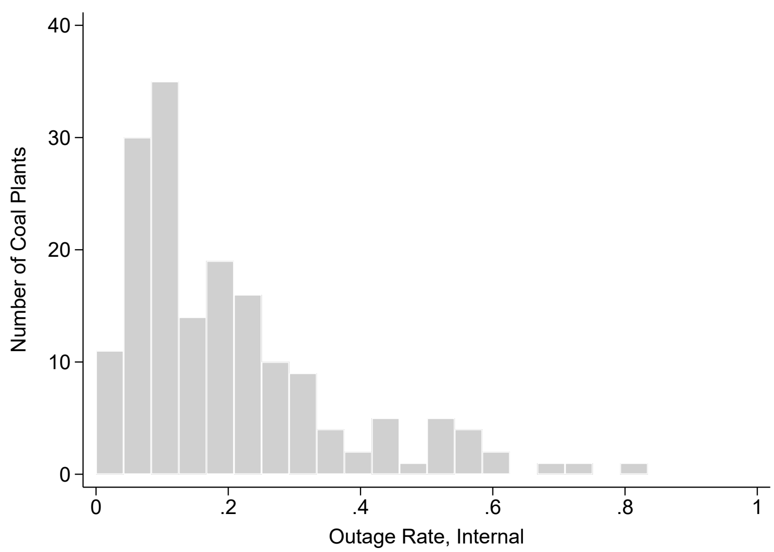

Notes: This figure reports histograms of rate at which Indian coal plants report outages. The variables on the horizontal axes divides each plant's total of MW-days on outage throughout our sample by the sum of its MW-days of capacity. By construction, this outage rate is between 0 and 1, and we drop days before (after) plants' first (last) observed generation before calculating outage rates. Each histogram contains 170 coal plants. Total outage rates include all outages in the numerator, classified as planned, force, or other. Internal outage rates include only outages "internal" to each facility (e.g., maintenance, but not transmission failures). An outage rate of 0.2 could mean 100 percent of capacity on outage during 20 percent of sample days, or 50 percent of capacity on outage during 40 percent of sample days.

Figure 9: Example expanded IEX supply curve

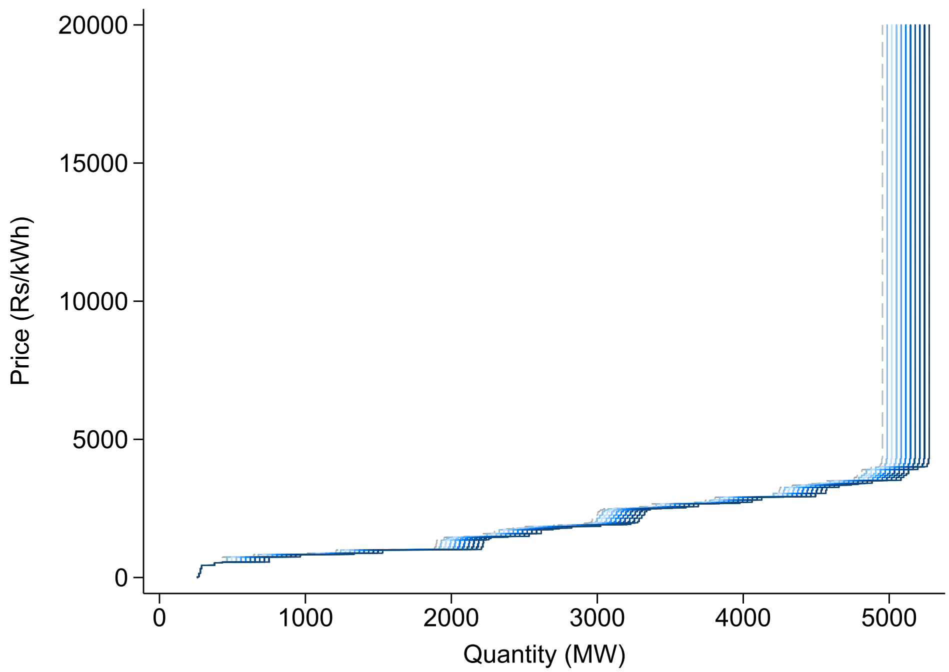

Notes: This figure provides an example of our exercise to increase the IEX supply curve. The dashed gray line shows the real data. The blue lines, from lightest to darkest, represent "counterfactual" supply curves in which we add in 10, 20, ..., up to 100 percent of capacity under outage for plants with a marginal cost well below the market clearing price. In this example, we consider only plants with a marginal cost of 50 percent of the IEX market clearing price or below.

Figure 10: Back-of-the-envelope calculations

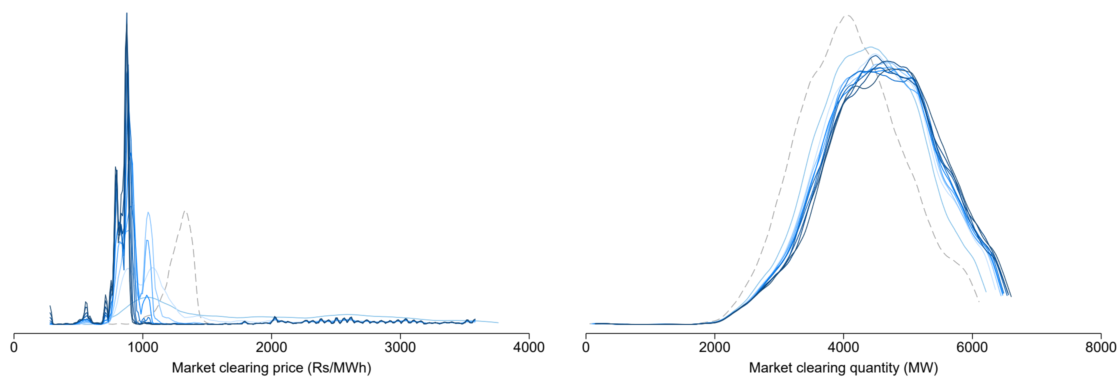

Notes: This figure displays the results of a back-of-the-envelope exercise in which we expand the supply of the IEX, drawing from plants under outage. In particular, we expand the supply curve by adding "unavailable" (a.k.a under outage) capacity from marginal plants – those with marginal costs below the IEX market clearing price – into the IEX. We conservatively only count plants whose marginal cost is below 50 percent of the market clearing price to be used. The left panel plots the actual (dashed gray) and counterfactual (blues) market clearing prices. The "counterfactual" prices range from the lightest blue, in which we return 10 percent of unavailable capacity to the market, to the darkest blue, in which we return 100 percent of unavailable capacity to the market. In the right panel, we plot market clearing quantities.