



Regular article

The health costs of dirty energy: Evidence from the capacity market in Colombia[☆]

Achyuta Adhvaryu^{a,b}, Teresa Molina^{c,d,*}, Anant Nyshadham^{a,b}, Jorge Tamayo^e, Nicolas Torres^f

^a University of California San Diego, NBER, J-PAL, BREAD, United States of America

^b Good Business Lab, United States of America

^c University of Hawaii at Manoa and UHERO, United States of America

^d IZA, Germany

^e Harvard Business School, United States of America

^f Good Business Lab, Colombia



ARTICLE INFO

JEL classification:

I18

O13

Q53

Keywords:

Electricity generation

Pollution

Health

Colombia

ABSTRACT

The health effects of “dirty” (fossil fuel driven) energy production are difficult to measure accurately due to the endogeneity of fuel choice. We exploit an electricity policy in Colombia that generates a price-based trigger for the use of thermal energy sources. Comparing municipalities near high versus low capacity plants, we first document that the activation of this trigger – which increased thermal energy production – led to significantly higher local pollution levels. This change increased cardiovascular-related ER mortality by 56% and respiratory-related morbidity by 9%. Our results translate to a cost of 996 million USD in terms of lives lost and higher healthcare costs.

1. Introduction

Governments around the world are faced with the choice of investing in clean or dirty energy. Global climate conferences often bring to light the complicated and controversial nature of these decisions (McGrath, 2021; Cursino and Falkner, 2021; de la Garza, 2021). It is clear that leaders perceive an inherent conflict between the reduction of dirty energy production and the promotion of other national interests (Rowlatt and Gerken, 2021; Geall, 2021; Hawkins, 2021).

This paper seeks to shed light on the tradeoff between clean and dirty energy by estimating the health costs of fossil fuel based energy production. This is a difficult task because the choice of fuel and the amount of power generated in a given region is typically endogenous, determined by a host of factors (including the preferences of the local population) which are also correlated with drivers of population health. Previous work has dealt with this endogeneity problem by exploiting exogenous shocks to power generation, including power plant closures,

expansions, and worker strikes (Ransom and Pope, 1995; Severini, 2017; Beach and Hanlon, 2018; Yang and Chou, 2018; Clay et al., 2021; Luechinger, 2014; Lavaine and Neidell, 2017). In this paper, we take advantage of a unique Colombian electricity pricing policy, in which an increase in thermal generation is triggered whenever the wholesale electricity price exceeds a pre-determined scarcity price. Our goal is to estimate how this ramp-up of thermal generation affects population health.

Though closely related to the large body of work documenting the negative effects of pollution on various health outcomes (e.g., Chay and Greenstone, 2003; Currie and Neidell, 2005; Jayachandran, 2009; Currie et al., 2014, and the studies cited in the previous paragraph), the research question of this study is distinct. Unlike these papers, which typically aim to recover the causal effect of a change in pollution levels, we are interested in the reduced form effect of fossil fuel based energy generation on health, which we argue is the policy relevant question of interest. A policymaker typically will have various policy

[☆] We thank Michael Roberts, Becca Taylor, and seminar participants at Michigan, GRIPS, UH Manoa, Notre Dame, Pac-Dev, and Monash for helpful feedback.

* Corresponding author.

E-mail addresses: adhvaryu@umich.edu (A. Adhvaryu), tmolina@hawaii.edu (T. Molina), nyshadha@umich.edu (A. Nyshadham), jtamayo@hbs.edu (J. Tamayo).

¹ This is a key distinction between this study and Ordoñez (2020), which aims to estimate the effect of PM 10 on health outcomes in the same setting. Also relying on the fact that thermal generation ramps up when hydropower is expensive, Ordoñez (2020) uses national river flows interacted with thermal power capacity as instruments for pollution levels. Because we are interested in identifying the effect of a policy lever, and for additional reasons described in Section 4.2, we choose to adopt a reduced form rather than an instrumental variables approach.

levers that can be used to switch from dirty to clean energy but will have less control over the amount of pollution actually emitted and, importantly, the exposure of the population to this pollution increase. Differences in the behavioral responses of individuals and the spatial distribution of a population will lead to different changes in pollution exposure in response to the same increase in pollutant emissions. From a government's perspective, therefore, the key question of interest is how health is affected by a policy that changes the generation fuel mix.¹ The possibility of mitigation and avoidance behavior, which may vary by socioeconomic status or other (unobservable) population characteristics, means that the policy parameter of interest is not easily recovered from estimates of the pollution effects of electricity generation and the health effects of pollution emissions.

Another important contribution of this study is its focus on Colombia: the vast majority of papers that exploit exogenous shocks to power generation has focused on the United States (Ransom and Pope, 1995; Severnini, 2017; Beach and Hanlon, 2018; Yang and Chou, 2018; Clay et al., 2021) or other high-income countries (Luechinger, 2014; Lavaine and Neidell, 2017). Recently, evidence from lower income countries has begun to emerge, focusing primarily on coal plants in India (Gupta and Spears, 2017; Barrows et al., 2021; Datt et al., 2021). Evidence from outside this setting is still very limited (Cesur et al., 2017; Ordoñez, 2020).

Estimates from lower income countries are important because the majority of the predicted increase in energy consumption is expected to come from non-OECD countries (US EIA, 2021), whose energy source choices will therefore be globally important. It is unclear how generalizable the evidence from rich countries will be. We might expect the health effects of fossil fuels in lower-income countries to be larger due to higher pollution levels (and potential non-linearity in the effects of pollution), lower health levels, and lower quality healthcare systems. On the other hand, fossil fuel generation may have less of an impact on health due to competing risks: there are other (potentially more important) drivers of mortality in lower income countries.

As mentioned above, we take advantage of an electricity pricing policy in Colombia, where the majority of electricity is generated by hydroelectric plants. On days when the wholesale electricity price exceeds a pre-determined level, thermal plants (which include coal, natural gas, diesel, and other liquid fuels) ramp up their generation. This typically happens because of very low rainfall restricting the supply of hydroelectricity.

Using daily data on electricity prices and generation, we are able to identify "scarcity days" as days when the wholesale price exceeds the scarcity price. Simply comparing health outcomes on scarcity and non-scarcity days would be unlikely to provide unbiased estimates of the health effects of thermal generation for two reasons. First, high wholesale prices are driven by demand and supply factors. If a scarcity day is triggered due to high demand for electricity, it would be difficult to separate the effects of increased thermal generation from the effects of the factors that drive electricity demand. A similar argument could be made for supply-side factors (in this case, primarily low rainfall), though we control flexibly for rainfall in our regressions. Second, the health data we use exhibits large day-to-day fluctuations in the extent of under-reporting, with particularly high under-reporting during the scarcity period (due to factors completely unrelated to electricity generation, as we discuss later). For these reasons, we make use of cross-sectional variation in addition to the scarcity day comparison to ensure that we are isolating the effect of the higher thermal generation that occurs on scarcity days.

Specifically, we characterize municipalities based on the average capacity of thermal power plants in their vicinity and categorize them into "high capacity" and "low capacity" municipalities by splitting at the median. Thermal plants with greater capacity are able to generate more electricity and therefore more pollution. This implies that municipalities near high capacity plants should be exposed to greater increases in pollution on a scarcity day, a hypothesis we are able to confirm

empirically.² That is, we regress various pollutant measurements on location fixed effects, time fixed effects, weather controls, and our main variable of interest: an interaction between a high capacity and scarcity day indicator. We document significantly larger increases in PM 2.5, PM 10, SO₂, and CO on scarcity days in high capacity compared to low capacity municipalities. Estimates of the interaction term correspond to a 36% increase relative to mean PM 2.5, 16% for PM 10, and 25% for SO₂.

Having documented that the interaction between high capacity and scarcity is a significant driver of pollution levels, we then use the same specification to estimate the effects of increased thermal generation on health outcomes. Existing work on the effect of power generation on health has almost exclusively focused on infant mortality and infant health as the outcomes of interest, but we are able to study a rich set of health outcomes. We have access to daily morbidity counts (specifically, the number of people who visited a health facility) by diagnosis code, as well as daily emergency room (ER) mortality counts by diagnosis code. We have data on all ages and can examine our outcomes (respiratory and cardiovascular morbidity and mortality) separately for infants, children, adults, and the elderly.

We find that respiratory morbidity increases significantly more for high capacity compared to low capacity municipalities on scarcity days; the magnitude of the interaction coefficient is 9% relative to the mean. This increase is accompanied by an increase in respiratory costs equal to 10% of the mean. We also find statistically significant and large effects on cardiovascular ER mortality, equivalent to 56% of the mean. These mortality effects are driven by the elderly.

Our main specification controls for municipality, date, and state-by-year fixed effects, along with flexible functions of temperature, precipitation and wind speed. In addition, we show that our results are not driven by geographic differences between high and low capacity municipalities (like altitude). Our results are also robust to allowing for different weather coefficients, month effects, and linear time trends across high and low municipalities. Back-of-the-envelope calculations reveal that, in terms of healthcare costs and lost lives, the cost of the scarcity period for the high capacity municipalities in our study was 996 million dollars (in 2015 USD).

2. Background

Colombia relies primarily on hydroelectric power, which generated over 70% of the country's electricity from 2000–2015 (McRae and Wolak, 2016). Almost all of the remaining electricity is generated by thermal power plants, which in Colombia's case include coal, natural gas, diesel, and other liquid fuels (see Appendix Figure A1 for the composition of total generation by technology for our sample period).

Colombia's dependence on hydropower can be problematic during times of low rainfall, as evidenced by the year of electricity rationing brought on by the El Niño event of 1992 (McRae and Wolak, 2019). Reforms that were largely motivated by this event eventually led to the development of the unique market structure and policy framework which provide us with the source of exogenous variation in fuel choice that we use to estimate the health costs of thermal energy.

Colombia's electricity market consists of a wholesale market (where wholesale electricity prices are determined daily), a retail market (where end users pay regulated prices for the electricity they consume), and a capacity market (where capacity payments made to generators are determined by auctions every few years). These capacity payments are paid to power plants even when they are not generating electricity. Generators that receive these payments are "obligated" to increase their generation whenever the wholesale market price exceeds a regulated "scarcity price". Specifically, whenever this happens, these generators

² This is consistent with the fact that plant capacity is highly correlated with excess capacity, as we show in Figure A4.

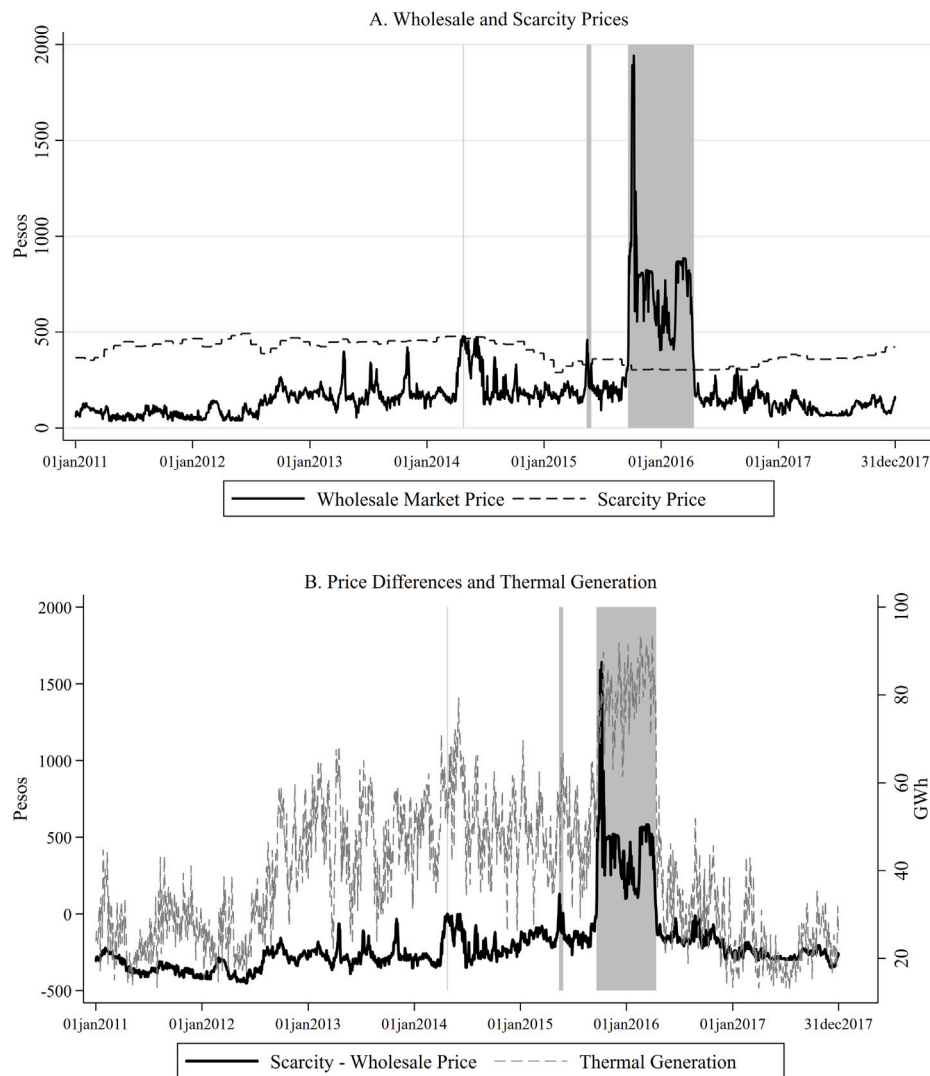


Fig. 1. Prices and Thermal Generation. Notes: Gray shaded areas denote scarcity days. Thermal generation is the total electricity generated across all thermal power plants.

must pay the difference between the wholesale price and the scarcity price, multiplied by their assigned generation capacity. This provides a financial incentive for generators to produce at least up to their assigned capacity, as they will end up charging the wholesale price, paying the difference between the scarcity price, and receiving the scarcity price (McRae and Wolak, 2019).

Panel A of Fig. 1 plots the daily wholesale market price (solid line) and scarcity price (dashed line) during our study period, 2011 to 2017. The gray shaded regions mark days on which the wholesale price exceeded the scarcity price, which we refer to as “scarcity days”. The primary scarcity period during these years took place between 2015 and 2016, caused by another El Niño event.

Panel B of Fig. 1 confirms that thermal plants do indeed increase generation during scarcity days. The gray dashed line, which represents total thermal electricity generation, jumps up during the scarcity period shaded in gray (when the difference between the wholesale and scarcity price exceeds zero). Appendix Figure A1 shows this increase is driven by several types of dirty energy (diesel, coal, and other liquid fuels), as well as natural gas, which is cleaner. While we would not expect large increases in pollution due to the increase in natural gas generation (which went from 19% of total generation prior to the scarcity period to 27% during the scarcity period), there could be substantial pollution effects driven by the increased diesel generation (which went from 1% to 9% of total generation), coal generation (9% to 10%) and generation

from other thermal sources (less than 1% to 3%). It is also clear from Appendix Figure A1 that generation from hydroelectric plants decreased during this period, indicating a shift away from renewable to thermal electricity sources during the scarcity period.

In this paper, we investigate what happens to pollution levels and, subsequently, health outcomes during these scarcity periods. Importantly, the increase in thermal generation that takes place on scarcity days is triggered by a pricing rule, rather than endogenous factors – like institutional quality, economic or political conditions, or technological improvements – that typically drive fuel choice decisions across countries and regions over time.

Motivated by the large body of work documenting links between air pollution and measures of respiratory and cardiovascular health specifically (Brunekreef and Holgate, 2002), we focus on these two disease categories in our analysis. Different pollutants affect health through different channels, but negative effects on respiratory health are generally driven by causing oxidative stress, inflammatory responses, and adverse changes in lung function (Kurt et al., 2016). Oxidative stress and inflammatory responses can also negatively affect the cardiovascular system, and some pollutants (PM 2.5) are fine enough to cross into the bloodstream, directly affecting the cardiovascular system (Brook et al., 2004, 2010).

Table 1
Summary statistics.

	All years	Pre-2015		
	All (1)	High capacity (2)	Low capacity (3)	Difference (4)
Thermal Generation (GWh)	2.40 (2.54)	3.31 (3.00)	1.40 (1.48)	1.91*** (0.12)
Respiratory morbidity (per 100,000)	22.79 (34.58)	22.68 (32.97)	23.62 (36.51)	-0.94 (1.34)
Cardiovascular morbidity (per 100,000)	33.35 (44.73)	31.90 (44.24)	30.35 (42.16)	1.55 (1.58)
Respiratory cost (per person)	9.29 (14.24)	9.14 (14.21)	9.59 (14.60)	-0.45 (0.60)
Cardiovascular cost (per person)	15.63 (26.06)	15.47 (26.19)	14.45 (24.99)	1.02 (0.96)
Respiratory ER mortality (per 100,000)	0.03 (0.55)	0.03 (0.37)	0.04 (0.47)	-0.011 (0.019)
Cardiovascular ER mortality (per 100,000)	0.01 (0.45)	0.01 (0.33)	0.01 (0.41)	-0.001 (0.0022)
Number of Health Facilities	110.51 (275.75)	107.28 (317.18)	105.09 (184.91)	2.19 (21.38)
Municipality Population	48,830 (350,882)	57,975 (460,453)	37,859 (151,198)	20,117 (28,704)
Municipality Share Children (0–14)	0.29 (0.04)	0.29 (0.04)	0.29 (0.04)	0.003 (0.003)
Municipality Share Prime-age Adults (15–64)	0.62 (0.03)	0.62 (0.03)	0.62 (0.03)	-0.003 (0.003)
Municipality Share Elderly (65 or more)	0.09 (0.03)	0.09 (0.03)	0.09 (0.03)	-0.001 (0.03)
Municipality GDP	642.82 (5962.20)	837.02 (8226.56)	445.86 (1698.45)	391.20 (498.20)
Municipality Educational Attainment	7.30 (1.10)	7.22 (1.05)	7.38 (1.14)	-0.16* (0.092)
Municipality Wind Speed (m/s)	2.32 (1.35)	2.66 (1.66)	2.41 (1.49)	0.25*** (0.06)
Municipality Rainfall (mm)	3.57 (5.28)	3.92 (5.64)	3.53 (5.16)	0.39*** (0.12)
Municipality Temperature (C)	11.23 (3.80)	12.41 (4.44)	11.14 (3.19)	1.27*** (0.24)
Municipality Altitude	1,418.04 (1,333.51)	1,255.23 (903.14)	1,578.01 (1,634.75)	-322.80*** (110.80)
Observations	1,449,819	417,846	410,541	828,387

Notes: Sample spans the years 2011–2017 and restricts to municipalities located within 100 kilometers of a thermal power plant. Unit of observation is a municipality-day.

3. Data

Drawing from several data sources, we construct a municipality-day-level panel spanning the years 2011 to 2017. We first restrict to municipalities that are close enough to a thermal power plant to be affected by changes in thermal generation. Using public information on power plant locations, we identify and restrict our main sample to municipalities located within 100 kilometers (calculated using the municipality's geographic centroid) of a thermal power plant. These municipalities are represented by the shaded regions in Appendix Figure A2.³ A 100 kilometer cutoff balances representativeness with the need to focus on municipalities that are close enough to be affected by a power plant. The resulting sample includes more than 70% of Colombia's population (567 municipalities). We also show robustness to a 120 kilometer cutoff, which includes 83% of the population.

3.1. Electricity generation

Our information about thermal power plants and electricity prices comes from the Colombian market operator XM. As mentioned above,

³ Because Colombia is divided by two large mountain ranges, a municipality that is physically close to a power plant may be very unlikely to be affected by it if it is on the opposite side of a mountain range. Therefore, when implementing the 100 kilometer cutoff, as with all cutoffs used in the remainder of the paper, we exclude any areas that are not in the same natural region (of which Colombia has six) as the point of interest.

we have daily spot prices and scarcity prices, which allow us to identify a scarcity day as any day when the spot price exceeds the scarcity price. Scarcity days account for 8% of the sample period.

This data source is what we use to split municipalities into two groups based on the capacity, or maximum generation potential, of their nearby power plants. Specifically, we calculate the inverse-distance weighted average capacity of power plants within 100 kilometers of each municipality, and split the sample at the median. 50.4% of municipalities are considered high capacity according to this definition.

In Table 1, we report summary statistics for the full sample in column 1 and compare high and low capacity municipalities in the remaining columns. Specifically, we restrict to years prior to 2015 (i.e., before the first major scarcity event took place), and report summary statistics for high capacity municipalities in column 2, low capacity municipalities in column 3, and the difference between the two groups in column 4. The first row of Table 1 reports average daily thermal power generation for each municipality. This is a weighted average of the electricity generated by all plants within 100 kilometers, weighting each value by the inverse of the distance between that thermal plant and the municipality (scaled so that weights sum to 1). As expected, average electricity generation is significantly higher in high capacity municipalities (more than double the generation of low capacity municipalities).

3.2. Health outcomes

We obtain morbidity and mortality measures from the Integrated Information System for Social Protection (SISPRO), which contains the

Individual Register of Health Services (RIPS). The RIPS collects detailed information about medical consultations, ER visits, hospitalizations, and medical procedures that take place in any Colombian health service institution. This allows us to calculate, for each municipality-day, the number of patients and total costs, broken down by the ICD-10 diagnosis code assigned to the visit.⁴ We use these ICD-10 codes to identify respiratory (J00–J99) and cardiovascular (I00–I99) conditions.

Although the RIPS data only captures illness among people who visit a health facility, we argue it is still a useful measure of population morbidity. Due to high insurance rates in Colombia, this measure captures a large share of people who are sick. According to Camacho and Mejía (2017), 70% of Demographic and Health Survey respondents who needed health treatments actually visited a health facility. Increases in our morbidity measures will be driven by increases in the number of people who are sick at all, as well as the share of people whose illness is severe enough for them to seek out formal healthcare.

Another feature to note about the RIPS is that there is likely to be substantial underreporting. As Appendix Figure A3 shows, there are large month-to-month fluctuations in the number of health facilities that report to the RIPS. These fluctuations are unlikely to be solely driven by fluctuations in the number of health facilities that receive any patients and likely represent some measurement error. Of particular concern is the sharp drop that occurs at the end of 2015, which is during the main scarcity period in our study. This drop coincides with (and was likely caused by) the liquidation of one of the public health insurance providers, which generated substantial chaos in the healthcare system (Ministerio de Salud y Protección Social, 2015; Barbosa and Monsalve S., 2017). This data issue is one reason why our empirical strategy relies on both cross-sectional and time variation, not just on a comparison between scarcity and normal days.

Table 1 reports daily morbidity for respiratory and cardiovascular diseases, measured as daily patient counts per 100,000 municipality residents. Morbidity, cost, and mortality outcomes are similar across high and low capacity municipalities prior to 2015. The number of health facilities reporting to the RIPS in each municipality (by month), which is included as a control variable in our later analysis, is also balanced across the two groups.

3.3. Municipality characteristics

We obtain other municipality-level characteristics from the National Administrative Department of Statistics (DANE), which we report in Table 1. Population size, age composition, and municipality GDP are similar across high and low capacity areas. Educational attainment is slightly higher for high capacity municipalities and this difference is statistically significant at the 10% level, though small in magnitude (amounting to about 2% of the mean).

We also have weather information from Colombia's Institute of Hydrology, Meteorology and Environmental Studies (IDEAM). This information contains daily measures of wind speed, rainfall, and temperature from 303 measurement stations. We assign weather variables to municipalities using inverse-distance weighting within a 100 kilometer radius.⁵ We also obtain average municipality altitude from Instituto Geográfico Agustín Codazzi (IGAC). Table 1 shows that high and low capacity municipalities do appear to have significantly different geographic characteristics, which could reflect systematically different location decisions of high and low capacity plants.

⁴ We are also able to further disaggregate by age, which we use in parts of our analysis.

⁵ When we examine pollution as an outcome variable, we assign weather variables to pollution stations using the same inverse-distance weighting procedure.

3.4. Pollution

We use information on pollution levels from Colombia's Air Quality Information Subsystem (SISAIRE). These data contain measures of PM 2.5, PM 10, SO₂, CO, NO₂, and O₃ by hour from 127 measurement stations. We report summary statistics for these pollution measures (at the station-day-level, for the entire 2011–2017 period) in Appendix Table A1.

4. Empirical strategy

Our goal is to measure the effects of thermal generation on municipality-level health. To do this, we use scarcity days – days on which the scarcity price exceeds the wholesale price – as a source of quasi-experimental variation. Scarcity days trigger increased electricity generation at thermal power plants (which in Colombia include dirty energy sources as well as natural gas). In this section, we first examine the relationship between scarcity days and pollution, across high and low capacity areas, and use these findings to motivate our empirical specification. We then describe the regression specifications for our main analysis, in which we estimate the effects of thermal generation on various health outcomes.

4.1. Pollution and scarcity days

We begin by investigating how switching to thermal generation affects pollution levels, using scarcity days a source of exogenous variation.⁶ Though scarcity days are defined by a rule-based trigger, simply comparing pollution levels on scarcity and non-scarcity days would likely fail to identify the causal effect of switching to thermal generation. Scarcity days tend to occur when rainfall is very low, for example, and controlling for precipitation could be an incomplete solution depending on the nature of the non-linearities in the relationship between pollution and weather. Scarcity days are also more likely to occur when the demand for electricity is high, which could be correlated with our outcomes of interest. In general, comparing scarcity to non-scarcity days would not allow us to control for any day-specific effects, which could be important if scarcity days coincide with other events that are correlated with pollution levels (or, for our later analysis, the quality of our health outcome data — for reasons described in Section 3.2).

We therefore leverage variation across space as well as over time. We exploit the fact that power plants with higher unused capacity will increase their electricity generation more on a scarcity day, compared to power plants with less unused capacity. Scarcity days should therefore result in larger increases in pollution in areas near a plant with high excess capacity. To test this, we estimate the following specification:

$$P_{st} = \delta_1 \text{High Capacity}_s \times \text{Scarcity}_t + \delta_2 X_{st} + \eta_s + \gamma_t + \epsilon_{st}. \quad (1)$$

P_{st} represents average pollution (either PM 2.5, PM 10, SO₂, CO, NO₂, or O₃) at measurement station s on day t and Scarcity_t is a scarcity day dummy variable. High Capacity_s is an indicator equal to 1 if the average capacity of power plants within 100 kilometers of station s is above the median. We create this indicator using total capacity and not unused capacity because unused capacity varies day-to-day (and is driven by potentially endogenous factors). Moreover, the correlation between total capacity and unused capacity is high (0.72), as we illustrate graphically in Figure A4. The vector X_{st} includes state-by-year fixed effects and cubic functions of rainfall, temperature, and

⁶ Note that “switching” is at the region level as opposed to the plant level — it refers to an increase in generation at thermal plants along with a reduction at hydroelectric plants, as opposed to a single plant switching from thermal to hydroelectric generation.

Table 2
The impact of thermal generation on pollution levels.

	(1) PM2.5	(2) PM10	(3) SO ₂	(4) CO	(5) NO ₂	(6) O ₃
Scarcity Day × High Capacity	7.22*** (1.37)	6.99** (3.22)	2.42** (1.00)	75.7 (61.3)	0.36 (1.96)	-1.63 (1.81)
Observations	26635	65129	30905	27950	25326	48810
Mean of DV	20.1	43.7	9.87	966.1	27.1	24.3

Notes: Standard errors (clustered at station level) in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions control for station fixed effects, date fixed effects, state-by-year fixed effects, and cubic functions of temperature, precipitation, and wind speed. Sample restricted to stations within 100 kilometers of a thermal power plant.

wind speed (generated as an inverse-distance weighted average of all weather stations within 100 kilometers of pollution station s).

We are interested in δ_1 , the coefficient on the interaction between the high capacity and scarcity day indicator. This captures the differential effect of a scarcity day in a high capacity compared to a low capacity area, which we interpret as the causal effect of switching to thermal generation. Because η_s controls for any location-specific unobservables and γ_t controls for any day-specific effects, the identifying assumption is that the difference in pollution levels between high and low capacity areas would have remained the same on scarcity days if thermal power plant generation had not been triggered.

Table 2 reports the regression results from Eq. (1), using PM 2.5, PM 10, SO₂, CO, NO₂, and O₃ as dependent variables. The interaction term is positive and significant for PM 2.5, PM 10, and SO₂, which means the increase in these pollutants on a scarcity day is significantly larger in high capacity areas. The estimates correspond to a 36% increase relative to mean PM 2.5, 16% increase relative to mean PM 10, and 25% increase relative to mean SO₂.

The interaction term (High Capacity _{j} × Scarcity _{t}) is a significant driver of changes in pollution that are large in magnitude. As we discuss in the next sub-section, we use this variable as the independent variable of interest in our main analysis.

4.2. Estimating the effects of thermal generation

We use the specification below to estimate the reduced-form effects of dirty energy on municipality-level morbidity, health costs, and mortality. Because the scarcity-by-high-capacity interaction drives changes in more than one type of pollutant, we use a reduced form approach instead of an instrumental variables strategy (where we would essentially have only one instrument for multiple endogenous variables). A reduced form approach is also preferred because there are only 127 pollution monitor locations, not evenly distributed across the 567 municipalities in our sample.

For municipality j on date t , we estimate

$$\sum_{k=t}^{t+2} Y_{jk} = \beta_1 \text{High Capacity}_j \times \text{Scarcity}_t + \beta_2 X_{jt} + \eta_j + \gamma_t + \epsilon_{jt}, \quad (2)$$

where Y_{jt} represents either morbidity rates, costs, or mortality rates (respiratory and cardiovascular). Like Deryugina et al. (2019), we use a three-day sum as our outcome variable to capture delayed effects (pollution on a given day affecting health the following day(s) instead of the same day) and to avoid picking up short-run displacement effects (pollution resulting in earlier health visits or deaths without actually increasing total counts). The vector X_{jt} includes state-by-year fixed effects, the number of health facilities reporting to the RIPS in municipality j in the month of time t , and cubic functions of rainfall, temperature, and wind speed (generated as inverse-distance weighted averages of all weather stations within 100 kilometers of municipality j).

The main coefficient of interest is β_1 , which we interpret as the effect of thermal electricity generation on Y_{jt} . Again, because of the inclusion of municipality (η_j) and day fixed effects (γ_t), identification comes from any differential changes in outcomes on scarcity days,

across high and low capacity municipalities. Table 1 shows that high capacity and low capacity municipalities are similar in terms of health outcomes, demographics, and socioeconomic status (in the years prior to the first scarcity period in our analysis). This provides support for our identifying assumption: that the gap in health outcomes across high and low capacity municipalities would have remained constant on scarcity days if thermal electricity generation had not been ramped up.⁷ We note, however, that high and low capacity municipalities do differ in terms of geographic characteristics (as shown in Table 1). We therefore run several robustness tests to ensure that these differences are not responsible for any differential trends in outcomes during the scarcity period.

5. Results

We begin by examining the effects of thermal electricity generation on morbidity, measured using the three-day total number of patients (per 100,000 municipality residents) categorized under a particular disease category. In column 1 of Table 3, there is a positive and significant coefficient on the interaction between scarcity day and high capacity. Switching to thermal generation increases the number of respiratory disease patients by 9% – approximately 6 additional patients per 100,000 residents. Column 2 reveals a positive but statistically insignificant coefficient for cardiovascular morbidity. Three-day total costs from respiratory disease increase by approximately 2.7 pesos per person (10% of the average cost) as a result of switching to thermal power generation (column 3). As with morbidity, the coefficient on cardiovascular costs is positive but statistically insignificant.

In addition to morbidity and costs, we also investigate whether thermal generation increases mortality. The RIPS data only record mortality for ER visits, which we use to calculate the three-day total of respiratory and cardiovascular ER deaths (per 100,000 people in a municipality). While there is no significant effect on respiratory mortality (column 5), column 6 shows that thermal generation increases cardiovascular ER mortality by 56% (0.023 deaths per 100,000 residents).

Table 4 explores heterogeneity by age. We define the following age categories: infants (under 1), children (between 1 and 14), youth and adults (between 15 and 64), and the elderly (65 or older). We calculate our morbidity and mortality outcomes for each age category and repeat our main regressions for each of these age groups.

Column 1 of Table 4 shows that those aged 15 and older are driving the effects on respiratory morbidity, with the coefficients for these groups representing about a 9%–10% increase relative to the mean. Effects on respiratory costs exhibit a similar age pattern. On the other hand, in column 6, the cardiovascular mortality effects are driven by the elderly, with an effect size of 87% relative to the mean.

⁷ Note that this is assumption would be less likely to hold if we used distance from thermal plant as our source of cross-sectional variation. Municipalities that are more than 100 km from a thermal plant are significantly different in terms of morbidity, mortality, and socioeconomic status, and could have responded differently to the chaos in the healthcare system coinciding with the 2015 scarcity period, which is why we do not use them as a control group in our analysis.

Table 3
The impact of thermal generation on morbidity, costs, and mortality.

	Morbidity per 100,000		Costs per person		Mortality per 100,000	
	Resp. (1)	Cardio. (2)	Resp. (3)	Cardio. (4)	Resp. (5)	Cardio. (6)
Scarcity Day × High Capacity	6.25** (2.77)	0.99 (3.52)	2.70*** (0.95)	2.64 (1.60)	0.025 (0.034)	0.023* (0.012)
Observations	1448685	1448685	1448685	1448685	1448685	1448685
Dep. Var. Mean	68.41	100.1	27.87	46.93	0.0976	0.0408

Notes: Standard errors (clustered at municipality level) in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All outcome variables are three-day totals. All regressions control for municipality fixed effects, date fixed effects, state-by-year fixed effects, number of health facilities reporting to the RIPS, and cubic functions of temperature, precipitation, and wind speed. Sample restricted to municipalities within 100 kilometers of a thermal power plant.

Table 4
The impact of thermal generation on morbidity, costs, and mortality, by age.

	Morbidity per 100,000		Costs per person		Mortality per 100,000	
	Resp. (1)	Cardio. (2)	Resp. (3)	Cardio. (4)	Resp. (5)	Cardio. (6)
A. Infants (Less than 1 year old)						
Scarcity Day × High Capacity	9.93 (10.5)	1.13 (0.90)	2.74 (2.46)	-0.084 (0.63)	-0.064 (0.40)	0.0086 (0.014)
Dep. Var. Mean	206.7	3.470	50.07	2.447	0.808	0.0138
B. Children (Ages 1–14)						
Scarcity Day × High Capacity	6.94 (4.48)	0.82** (0.32)	1.56 (1.06)	0.25 (0.16)	0.063 (0.066)	-0.0010 (0.0012)
Dep. Var. Mean	96.39	2.373	30.26	1.286	0.147	0.00151
C. Youth/Adults (Ages 15–59)						
Scarcity Day × High Capacity	4.13** (1.82)	0.26 (1.74)	0.94** (0.44)	1.18 (0.97)	0.014 (0.018)	-0.0023 (0.0058)
Dep. Var. Mean	41.01	54.56	13.22	26.41	0.0421	0.0157
D. Elderly (Over 60 years old)						
Scarcity Day × High Capacity	10.8** (5.31)	-0.32 (19.3)	3.71** (1.68)	10.8 (10.00)	0.050 (0.047)	0.20** (0.079)
Dep. Var. Mean	120.3	536.2	44.84	260.5	0.180	0.230
Observations	1448685	1448685	1448685	1448685	1448685	1448685

Notes: Standard errors (clustered at municipality level) in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All outcome variables are three-day totals. All regressions control for municipality fixed effects, date fixed effects, state-by-year fixed effects, number of health facilities reporting to the RIPS, and cubic functions of temperature, precipitation, and wind speed.

We explore how these effects vary by municipality-level socioeconomic status in Table A2. To proxy for socioeconomic status, we calculate the average education level for each municipality using data from 2011. We repeat our analysis separately for municipalities with average education below and above the municipality-level median. Coefficient estimates are larger for the low education group in the respiratory morbidity and cardiovascular mortality regressions, while the opposite is true for the remaining outcome variables. The differences between the groups are not statistically significant.

5.1. Robustness checks

Our main identifying assumption is that high and low municipalities would have seen similar changes in pollution and health during the scarcity period if the increase in thermal generation had not been triggered. One concern is that high and low municipalities differ in geographic characteristics like altitude, rainfall, and temperature (as shown in Table 1), and it could be the case that these underlying differences generated diverging trends in outcomes during the El Niño event. To account for this possibility, we add interactions between the scarcity day dummy and municipality altitude, average temperature (pre-2015), and average rainfall (pre-2015). The results reported in panel A of Table A3 reveal that our results are robust to the inclusion of these controls. Next, in panel B of Table A3, we allow for weather variables to have different effects in high and low capacity municipalities. This helps ensure our coefficient estimates are not being driven by differential responses to the El Niño event (specifically, the accompanying changes in weather) responsible for the scarcity period in our study. In panel B, we allow for different seasonal trends for high and low capacity

municipalities (by controlling for group-specific month fixed effects). In panel C, we allow for different quadratic trends for high capacity and low capacity municipalities. None of these additional controls substantially alter coefficient estimates.

Next, we conduct a falsification test, using morbidity, costs, and mortality from external causes (ICD-10 codes V00-Y99, which include accidents) as our outcomes of interest. If our results above were driven by changes in health-seeking behavior as opposed to changes in health levels, we would expect to see a significant coefficient on our interaction term of interest in these regressions. Appendix Table A4 reveals no significant effects of the scarcity by high capacity interaction, suggesting this is not the case.

We also examine whether there are any significant changes in the gap between high and low capacity municipalities during the week before and after a scarcity period. Table A5 repeats our original regression and adds two additional interaction terms: high capacity interacted with an indicator for the week before a scarcity period, and high capacity interacted with an indicator for the week after a scarcity period. The former should yield statistically insignificant coefficients if it is indeed the pollution generated on scarcity days that is driving our effects. The latter will reveal any persistent effects of the pollution increases.

Across all columns of Table A5, the scarcity day interactions are similar in magnitude to our baseline estimates, and the interactions with the week-before indicator are all small and statistically insignificant, providing further support for the validity of our empirical strategy. The coefficients on the week-after interaction term are all larger in magnitude. In the regression on cardiovascular mortality, it is statistically significant and even larger in magnitude than the scarcity

day interaction term (column 6). The pollution increases on scarcity days appear to continue to affect cardiovascular mortality even after the scarcity period is over, perhaps indicating that it takes some time (and perhaps continued exposure) for the health effects of increased pollution to translate into higher mortality.

We also run event study regressions for outcomes significantly affected by the switch to thermal generation (respiratory morbidity, respiratory costs, and cardiovascular mortality). Summarized in Figure A5, these regressions yield similar conclusions. The outcome variables are single-day morbidity or mortality counts (per 100,000 residents). The main regressors of interest are interactions between indicators for every quarter and the high capacity interaction (leaving the quarter just before the main scarcity period as our omitted category). We report the results for all outcomes for which significant coefficients were reported in Table 3. Red dots represent quarters in which a scarcity day took place, while blue crosses represent all other months.

For respiratory morbidity (in panel A), the blue crosses display a relatively flat pattern prior to the first scarcity day; the majority of coefficients are small and statistically insignificant. On the other hand, the red dots are positive and statistically significant throughout most of the 2015–2016 scarcity period. Interestingly, coefficients remain positive (and in most cases statistically significant) until the end of 2016, suggesting some persistence in the effects of the pollution increases during the scarcity period. The results for respiratory costs (panel B) show similar patterns: positive and significant coefficients during the scarcity period.

The results for cardiovascular mortality are less precisely estimated. Like in the other panels, the blue coefficients reveal no pre-trends. The first red scarcity dot (in quarter 2 of 2015) is higher than the previous period, though it is not statistically significant. Similarly, two of the three quarters of the longer scarcity period are positive (but insignificant). The last coefficient the largest in the whole series.⁸ This suggests that while the morbidity effects of pollution may be immediate, prolonged exposure may have been what caused the increases in mortality.

Finally, we demonstrate the robustness of our results to a 120 kilometer cutoff (Tables A6 and A7).

5.2. Back-of-the-envelope calculations

Using our coefficients estimated above, we use back-of-the-envelope calculations to estimate the cost of the scarcity period in terms of increased healthcare costs and lost lives, for high capacity municipalities. First, the interaction coefficient of 2.7 pesos per person per municipality per day in the respiratory cost regression (column 3 of Table 3) translates into an increase of 3.5 million USD (in 2015 dollars) for high capacity municipalities throughout the entire scarcity period.⁹

To calculate mortality costs, we use the ER cardiovascular mortality coefficient (0.023 deaths per 100,000 people per municipality per day) and the value of a statistical life calculated specifically for Colombia by Viscusi and Masterman (2017): 1.228 million 2015 USD. This yields an estimate of 992.7 million USD (for high capacity municipalities throughout the entire scarcity period), much larger than the total costs stemming from increased healthcare utilization. Therefore, our results translate to a cost of 996 million USD for high capacity municipalities, which is a conservative estimate of the total cost of the policy because it ignores any costs experienced by low capacity municipalities.

⁸ Note that this month includes non-scarcity days at the end of the quarter.

⁹ This calculation uses an exchange rate of 2745 pesos per 2015 USD, an average high capacity municipality population of 57,970, and sums across 212 scarcity days and 286 high capacity municipalities.

6. Conclusion

This paper takes advantage of a unique electricity policy in Colombia to obtain causal estimates of the health costs of switching to thermal energy generation. Comparing municipalities near high capacity plants to those near low capacity plants, on days when a price trigger substantially increases thermal generation, we find that PM 2.5, PM 10, and SO₂ levels increase significantly more in high capacity municipalities.

Using this same regression specification, we estimate the effects of increased thermal generation on morbidity and mortality outcomes. Thermal generation increases respiratory morbidity (primarily for those older than 15) and cardiovascular mortality (primarily for the elderly). We calculate that, for high capacity municipalities, the entire scarcity period led to 996 million USD worth of healthcare costs and lost lives.

CRedit authorship contribution statement

Achyuta Adhvaryu: Conceptualization, Methodology, Review & editing. **Teresa Molina:** Methodology, formal analysis, Writing. **Anant Nyshadham:** Conceptualization, Methodology, Review & editing. **Jorge Tamayo:** Conceptualization, Data curation, Methodology. **Nicolas Torres:** Data curation, Formal analysis, Writing.

Data availability

The authors do not have permission to share data.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jdeveco.2023.103116>.

References

- Barbosa, C.V., Monsalve S., M.M., 2017. El origen de la enfermedad de Cafesalud. In: El Espectador.
- Barrows, G., Garg, T., Jha, A., 2021. The health costs of coal-fired power plants in India. SSRN Electron. J.
- Beach, B., Hanlon, W.W., 2018. Coal smoke and mortality in an early industrial economy. *Econom. J.* 128 (615), 2652–2675.
- Brook, R.D., Franklin, B., Cascio, W., Hong, Y., Howard, G., Lipsett, M., Luepker, R., Mittleman, M., Samet, J., Smith, Jr., S.C., et al., 2004. Air pollution and cardiovascular disease: a statement for healthcare professionals from the Expert Panel on Population and Prevention Science of the American Heart Association. *Circulation* 109 (21), 2655–2671.
- Brook, R.D., Rajagopalan, S., Pope, III, C.A., Brook, J.R., Bhatnagar, A., Diez-Roux, A.V., Holguin, F., Hong, Y., Luepker, R.V., Mittleman, M.A., et al., 2010. Particulate matter air pollution and cardiovascular disease: an update to the scientific statement from the American Heart Association. *Circulation* 121 (21), 2331–2378.
- Brunekreef, B., Holgate, S.T., 2002. Air pollution and health. *Lancet* 360 (9341), 1233–1242.
- Camacho, A., Mejía, D., 2017. The health consequences of aerial spraying illicit crops: The case of Colombia. *J. Health Econ.* 54, 147–160.
- Cesur, R., Tekin, E., Ulker, A., 2017. Air pollution and infant mortality: Evidence from the expansion of natural gas infrastructure. *Econom. J.* 127 (600), 330–362.
- Chay, K.Y., Greenstone, M., 2003. The impact of air pollution on infant mortality: evidence from geographic variation in pollution shocks induced by a recession. *Q. J. Econ.* 118 (3), 1121–1167.
- Clay, K.B., Lewis, J., Severnini, E., 2021. Canary in a coal mine: Infant mortality, property values, and tradeoffs associated with mid-20th century air pollution. SSRN Electron. J.
- Currie, J., Neidell, M., 2005. Air pollution and infant health: what can we learn from California's recent experience? *Q. J. Econ.* 120 (3), 1003–1030.
- Currie, J., Zivin, J.G., Mullins, J., Neidell, M., 2014. What do we know about short-and long-term effects of early-life exposure to pollution? *Annu. Rev. Resour. Econ.* 6 (1), 217–247.
- Cursino, M., Falkner, D., 2021. COP26: China and India must explain themselves, says Sharma. BBC News, <https://www.bbc.com/news/uk-59280241>. (Online accessed 16 December 2021).
- Datt, G., Maitra, P., Menon, N., Ray, R., Dey, S., Chowdhury, S., 2021. Coal Plants, Air Pollution and Anaemia: Evidence from India. *J. Dev. Stud.* 59 (4), 533–551.
- de la Garza, A., 2021. Fossil Fuel Companies Are Still Influencing COP26, Despite Losing Their Official Role. TIME, <https://time.com/6110667/fossil-fuel-companies-cop26-influence/>. (Online accessed 1 April 2022).

- Deryugina, T., Heutel, G., Miller, N.H., Molitor, D., Reif, J., 2019. The mortality and medical costs of air pollution: Evidence from changes in wind direction. *Amer. Econ. Rev.* 109 (12), 4178–4219.
- Geall, S., 2021. China, Climate Politics and COP26. JSTOR, (Online accessed 1 April 2022).
- Gupta, A., Spears, D., 2017. Health externalities of India's expansion of coal plants: Evidence from a national panel of 40,000 households. *J. Environ. Econ. Manag.* 86, 262–276.
- Hawkins, W.R., 2021. At COP26, climate action remains hostage to national interests. *The National Interest*, <https://nationalinterest.org/feature/cop26-climate-action-remains-hostage-national-interests-196755>. (Online accessed 1 April 2022).
- Jayachandran, S., 2009. Air quality and early-life mortality evidence from Indonesia's wildfires. *J. Hum. Resour.* 44 (4), 916–954.
- Kurt, O.K., Zhang, J., Pinkerton, K.E., 2016. Pulmonary health effects of air pollution. *Curr. Opin. Pulm. Med.* 22 (2), 138.
- Lavaine, E., Neidell, M., 2017. Energy production and health externalities: Evidence from oil refinery strikes in France. *J. Assoc. Environ. Resour. Econ.* 4 (2), 447–477.
- Luechinger, S., 2014. Air pollution and infant mortality: A natural experiment from power plant desulfurization. *J. Health Econ.* 37 (1), 219–231.
- McGrath, M., 2021. COP26 climate summit president says progress made, but not enough. *BBC News*, <https://www.bbc.com/news/science-environment-57975025>. (Online accessed 16 December 2021).
- McRae, S.D., Wolak, F.A., 2016. Diagnosing the Causes of the Recent El Nino Event and Recommendations for Reform. Program on Energy and Sustainable Development, Stanford University.
- McRae, S.D., Wolak, F.A., 2019. Market Power and Incentive-based Capacity Payment Mechanisms. Working Paper, pp. 1–44.
- Ministerio de Salud y Protección Social, 2015. Gobierno inicia proceso de liquidación de saludcoop EPS. <https://www.minsalud.gov.co/Paginas/Gobierno-inicia-proceso-de-liquidaci'on-de-Saludcoop-EPS.aspx>. (Online accessed 30 November 2021).
- Ordoñez, P.J., 2020. Power Plants, Air Pollution, and Health in Colombia. Working Paper.
- Ransom, M.R., Pope, A.C., 1995. External health costs of a steel mill. *Contemp. Econ. Policy* 13 (2), 86–97.
- Rowlatt, J., Gerken, T., 2021. COP26: Document leak reveals nations lobbying to change key climate report. *BBC News*, <https://www.bbc.com/news/science-environment-58982445>. (Online accessed 16 December 2021).
- Severnini, E., 2017. Impacts of nuclear plant shutdown on coal-fired power generation and infant health in the Tennessee Valley in the 1980s. *Nat. Energy* 2 (4).
- US EIA, 2021. International energy outlook 2021.
- Viscusi, W.K., Masterman, C.J., 2017. Income elasticities and global values of a statistical life. *J. Benefit-Cost Anal.* 8 (2), 226–250.
- Yang, M., Chou, S.Y., 2018. The impact of environmental regulation on fetal health: Evidence from the shutdown of a coal-fired power plant located upwind of New Jersey. *J. Environ. Econ. Manag.* 89, 94–115.