The Coal Transition and Its Implications for Health and Housing Values

Rebecca Fraenkel,[*] Josh Graff-Zivin[†] and Sam Krumholz[‡]

Abstract

From 2005–2020, one-third of US coal plants had at least one coal-fired generator close. We utilize this natural experiment to estimate the effect of coal plant exposure on mortality and house values. Using a difference-in-differences design, we find that counties within 30 miles of a closing unit experience large health effects following shutdown. While these health improvements appear to capitalize into housing values, they only do so within 15 miles of the plant and only when the retirement is of all units. Taken together, these results underscore the importance of risk salience in shaping market-mediated price effects.

Keywords: coal, air pollution, power plant, energy generation, housing capitalization

JEL Classification: I18; Q40; Q52

1 Introduction

The notion that the value of local public goods should capitalize into housing values can be

[*] University of California, San Diego, rfraenkel@ucsd.edu

[†] University of California, San Diego, jgraffzivin@ucsd.edu

[‡] US Department of Justice, Antitrust Division krumholz.samuel@gmail.com; The author does not purport that his views reflect those of the Department of Justice.

doi:10.3368/le.100.1.102222-0092R

traced back to the seminal work of Rosen, 1974. This hedonic pricing approach has subsequently been used to estimate the value of many (quasi) non-market goods. These estimates are particularly important in assessing the value of environmental quality since benefits estimation is needed to justify regulatory actions (Congressional Research Service, 2017). The empirical literature on this subject spans a wide range of pollutants that include solid waste, criteria air pollutants, and air toxins to name a few (e.g. Chay and Greenstone, 2005; Greenstone and Gallagher, 2008; Haninger, Ma, and Timmins, 2017; and Currie et al., 2015). It is also worth noting that this corpus of research deploys a varied set of econometric approaches, some of which are more susceptible to bias due than others (see Bishop et al., 2020 for a good overview).

In this paper, we build upon this literature by exploiting transaction-level housing data to estimate the impacts of changes in air quality resulting from a recent and pervasive phenomenon – the retirement of coal-fired power plants. Over the past 15 years, the United States has witnessed an unprecedented decline in coal-fired energy production. Since 2005, total coal generation has fallen by 25% and more than 30% of plants have seen at least one generating unit retire.[1] At the same time, transmission constraints and other factors often mean that new energy generation sources, most often in the form of natural gas plants, are built nearby to closing ones, creating alternative sources of pollution (Burney, 2020). Indeed, in our sample, we see that on average counties near closing coal units experience little change in generation levels, but a large shift in generation mix. Thus, the health, employment, and blight effects from closures and the degree to which they manifest in housing values should reflect the net effect from these changes.

Our approach builds upon the seminal work of Burney (2020), who showed that closures led to reductions in pollution and mortality in the county that housed that power plant. We exploit the fact that pollution travels considerable distances from the plant (Changotra et al., 2021), to implement a difference-in-differences design, which compares changes in outcomes for individuals that vary in their proximity to a closing coal-fired unit. This allows us to examine the differential impacts of closures on pollution, health and employment. Since housing prices should capitalize both health and employment effects, we can then explore whether they do so at the same rate across geographies to see whether the salience of a closure, which is presumably higher for homeowners closer to the plant, influences valuations.

We begin our analysis with a focus on health and find that counties near a closing unit experience large mortality improvements following shutdown. Specifically, relative to counties between 30 and 45 miles away from a shutdown, counties whose population centroids are within 15 miles of a shutdown experience a 2.77% decline in cardiovascular mortality, while those that are 15–30 miles away experience a 1.47% decline. Counties within 15 miles of a shutdown also experience a small (albeit insignificant) decrease in respiratory mortality. Consistent with the etiology of disease resulting from power-plant pollution, neither group of counties experiences differential changes in non-cardiovascular, non-respiratory mortality after unit closure. These results are robust to the inclusion of additional covariates, including pre-closure generation levels, demographics, and mortality rates.

Next, we examine local economic impacts at the county level. We find no evidence for economically meaningful effects on either major economic outcomes or local mining in the two years following plant closure, regardless of distance from the plant, suggesting that the local economic effects of power plant shutdowns are negligible. This finding provides additional support that the observed mortality changes are caused by changes in pollution and not other social or economic factors caused by power plant shutdowns. It also suggests that our housing estimates can, with an appropriate amount of introspection, be viewed as a measure of willingness to pay for health.

Turning to our housing analysis, we find that housing values begin to increase within 6-10 months after retirement, but those effects are limited to those houses within 15 miles of the closing unit. Moreover, that effect is only significant for complete plant retirements (as opposed to partial ones), yielding an approximate increase in housing values of 5%. That there are no capitalization effects beyond 15 miles, despite the notable health effects at that distance, along with the absence of effects even quite close to the plant when the closure is partial, underscores the importance of subjective perceptions in shaping market-mediated price effects. These sorts of informational market failure have potentially far-reaching implications for the broader literature focused on pollution and housing values and the welfare implications that are often derived from them.

The rest of this paper proceeds as follows. Section 2 provides background on the size and speed of the US energy transition and the pollutants associated with this energy generation. Section 3 describes our data and empirical strategy. Section 4 presents our results and Section 5 concludes.

2 Background

US Energy Transition

Since 2005, US coal-fired generation has fallen by roughly 25% and more than 30% of all coal plants have had at least one unit come offline.[2] While the causes of the decline in coal generation and pollution are multi-faceted (e.g. competition from natural gas (Johnsen, LaRiviere, and Wolff, 2017), energy demand shocks (Linn and McCormack, 2017), competition from renewables (Fell and Kaffine, 2017) and increasingly strict regulatory policies (Krumholz, 2019)), recent papers have suggested that retirements can be almost entirely explained by market forces rather than stricter regulatory polices (Linn and McCormack, 2017; Davis, Holladay, and Sims, 2021).

The top left figure of Appendix Figure A.2 shows the cumulative proportion of Acid Rain Program coal plants operating in 2004 with at least one unit retired by year. By the end of the period, around 30% of plants have had at least one unit retire. The remaining plots show the probability of retirement as a function of 2004 unit SO_2 pollution rate, unit age, and average monthly 2004 unit generation. Older, smaller and more polluting plants are much more likely to have a unit retire. These findings motivate our identification strategy: given the significant differences in retiring and non-retiring plants, we restrict our analysis to retiring plants only.

Pollutants and health

Coal-fired power generation may adversely affect human health through a number of channels.

First, coal-fired power plants emit high quantities of SO_2 and NO_X. These molecules act as precursors to PM 2.5, which has been found to negatively affect human health across a wide variety of epidemiological and economic settings (EPA, 2018). In particular, PM 2.5 has been found to aggravate cardiovascular and respiratory disease, increasing the risk of heart attacks and respiratory failure (Dominici et al., 2006). Some epidemiological studies have also found that sulfur dioxide alone can have independent effects on cardiovascular and respiratory mortality (Katsouyanni et al., 1997), while NO_X acts as a precursor to ozone, which has been found to lead to increased respiratory and cardiovascular mortality (Bell et al., 2004).

In addition to the increase in particulate matter created by plant emissions, transporting coal to existing plants can also be a substantial source of PM 2.5 through emissions from trucks and trains. Coal plants' storage of fuel stocks and waste can also adversely affect human health. For instance, Jha and Muller (2017) find that increases in coal stockpiles increase local PM 2.5 levels, leading to higher rates of adult and infant mortality.

At the same time, power plants fueled by natural gas, while less polluting than coal plants, also emit significant quantities of PM 2.5 and are not health-neutral (Brewer et al., 2016; Burney, 2020). Since many of these coal-fired units are replaced with natural gas units (a claim that we will confirm directly with our data), the health effects of coal unit retirements should be viewed as the net improvements associated with the transition from coal to gas. Since the primary channel through which health is harmed is either respiratory or cardiovascular impairment regardless of the fuel source, our primary health analyses in this paper will look at the effects of closures on respiratory and cardiovascular mortality, while treating all other causes of death as a placebo outcome.

3 Data and Empirical Strategy

Data

Data for this project comes from several primary sources. Mortality data come from the 2005–2015 National Center for Health Statistics' restricted-use mortality file. These data contain detailed information on all deaths in the United States within this time window including month and year of death, cause of death and the deceased's home county. We use these data to create cause-specific mortality rates for each county for every month-year.[3] Specifically, we examine three causes of death: cardiovascular, respiratory and all other (total−cardiovascular−respiratory). Together, cardiovascular and respiratory deaths make up approximately 42% of all deaths.

Data on unit fuel type, pollution control adoption, generation and pollution levels comes from the EPA's Continuous Emissions Monitoring System (CEMS). We use all plants covered by the Acid Rain Program, which covers the vast majority of US coal units with nameplate capacity greater than 25 megawatts.

Data on generator retirements comes from EIA form 860, which lists retirement dates for each retiring generator. We restrict our sample of retiring units to only those with a 2005 nameplate capacity above 25 MW. 415 coal plants were covered by the acid rain program and used coal to generate electricity in 2004. Among these 138 had at least one unit that retired or planned to retire before 2018. A further 33 plants had retiring coal units, but were uncovered by the Acid Rain Program and so were excluded from our sample. We finally dropped all plants with no coal-fired production in 2004, the year before our sample period begins, resulting in an additional 33 plants being excluded. Ultimately, we examine the effects of closure on 383 plants spread across 40 states of which 134 had at least one unit that retired before 2018.

Data on annual county demographic and economic characteristics comes from the US Census Bureau intercensal estimates of age and population, the Small Area Income and Poverty Estimates (SAIPE) and the Quarterly Census of Employment and Wages (QCEW).

Data on home sales come from Zillow-ZTRAX. This dataset contains detailed information about home sales and characteristics and includes precise coordinates that allow us to better understand differences in the effects of pollution across relatively short distances. A detailed description of how we assemble the housing data can be found in Appendix B. Some plants cannot be included in our analyses because there are too few housing transactions near them or because they are in non-disclosure states.[4] We link home sales data to plant retirements using the geographic coordinates of homes and plants. We use these coordinates to measure the distance from plants to homes and associate sales with plant retirements that are within 45 miles. If a transaction is in the analysis time window and is exposed to retirement of more than one plant, we associate the transaction only with the first retirement.

3.1 Empirical Strategy

The goal of this paper is to estimate the mortality and home price effects of coal unit closures. To do this, we need to construct a valid counterfactual for areas surrounding a coal unit closure subsequent to the unit closing. Using either areas surrounding non-closing plants or areas without coal plants is problematic as these counties differ from closing counties on a number of important dimensions. As an alternative strategy, we instead perform an event study design using only counties surrounding coal plants that closed (or are planning to close). To further control for area differences that may be correlated with the timing of closings, we compare counties closer to the closing with those that are located further away. Following Clay, Lewis, and Severnini (2016), we compare counties within and beyond 30 miles from a closing unit as we expect most of the direct health effects of closure to be within this radius.

Specifically, for each county we identify the first retirement that is within a 45 mile radius of the county's population centroid.[5] We then compare mortality changes in counties that are within 15 miles of the closing unit and within 15–30 miles of the closing unit to counties that are 30–45 miles from the closing unit. Accordingly, we estimate:

$$Y_{ct} = \gamma Post_{ct} + \sum_{b=1}^{D} \beta_b Post_{ct} * I_c(dist = b) + \alpha_c + \delta_{d(c)t} + \epsilon_{ct}. \quad (3.1)$$

where c indexes counties, $d(c)$ indexes Census divisions (which contain counties), t indexes month by year and b indexes distance bins where $b \in ([0,15), [15,30), [30,45))$. This specification non-parametrically estimates the effect of being various distances from the closing coal unit relative to the omitted distance bin (30–45 miles). If we assume conservatively that the coal unit has no mortality effect on counties whose population centroid is more than 30 miles away from the closing unit, the variables β_b describe the mortality effect of being in distance bin b. The inclusion of month by year by Census division effects implies that we are controlling for any time-variant trends in our outcomes within a Census division. We cluster standard errors at the plant level.

The identifying assumption of these approaches is that there is no unobserved factor correlated with both mortality and a coal unit's closing that disproportionately affects counties close to the closing plant following closure. We attempt to empirically check this assumption in several ways. First, we examine whether counties different distances from the closing units have differential mortality trends prior to closing. This tests whether farther counties serve as a valid control group for those closer to closing plants. Under this assumption, we estimate the following:

$$Y_{ct} = \sum_{m=-T}^{T} \gamma_m I_{ct}(t-\tau=m) + \sum_{m=-T}^{T} \sum_{b=1}^{D} \beta_{bm} I_{ct}(t-\tau=m) * I_c(dist = b)) + \alpha_c + \delta_{d(c)td(c)t} + \epsilon_{ct} \quad (3.2)$$

where m indexes months relative to closure and τ is time of closure.[6]

Second, if observed mortality differences are the result of coal unit closures, we should see closure have little effect on non-respiratory, non-cardiovascular mortality outcomes.

Accordingly, we use non-respiratory, non-cardiovascular mortality as a placebo outcome to examine if areas differentially close to the closure are experiencing shifts in overall mortality that may be unrelated to the changes in coal generation. Third, if results are not being driven by baseline differences across counties different distances away from the closing units, then coefficients should not change when we introduce baseline covariates interacted with time as covariates. We include these controls as robustness checks for each of our primary models. Finally, if this assumption is valid, then we should not see differential changes in other economic outcomes following a unit shutdown for counties close to the closing unit. We check this empirically using changes in employment, wages, population and income as outcome variables.

Appendix Figure A.1 shows the counties included in our sample broken down by distance bin. Our analysis is focused on the Eastern half of the United States as this is where the bulk of the country's coal generation is located. We also compare baseline characteristics of counties by their distance from a closing unit in Table A.1. Counties are largely comparable across distance bins, though mortality and poverty are marginally lower in the middle distance bin. Since we see the largest mortality improvements in the closest counties, and we do not see changes in non-pollution related deaths following plant closure, we are reassured that our results are from the closures themselves and not due to mortality trends in the surrounding communities unrelated to plant closures.

Closures and Home Prices

We examine the effect of plant closures on home prices using both a repeat sales framework and a cross-sectional hedonic regression.

$$Ln(Price)_{isdpt} = \beta_b I_{bip} Post_{pt} + \alpha_{bis} + \tau_{pt} + D_{dt} + \varepsilon_{ispt} \tag{3.3}$$

We examine the log price of home i in 2014 dollars, during year-month t, located in distance bin b, state s, and school district d, associated with plant p. Our primary specification is a repeat sales analysis which replaces home characteristic controls with an indicator variable for each home-plant pairing. In the hedonic regressions, we control for home characteristics by state and distance bin by plant. Home characteristics include binned number of bedrooms, bathrooms, square footage, lot size, year built, property type, roof type, and heat type. All categories include an indicator for missing values. The vector τ_{pt} contains plant by year by month fixed effects to control for any time-varying characteristics of homes and markets surrounding plants. D_{dt} controls for district by year effects to account for possible changes in tax revenue from plant closures that may pass through to home prices. (Fraenkel and Krumholz, 2022) The coefficient β on $Post_{pt} * I_{ip}$ is our coefficient of interest and describes how home prices change in the period following plant closure for homes located within distance bins of $([0,15), [15,30), [30,45))$. We measure the exact distance from each sold home to each plant and bin transactions accordingly. These distance bins match the distance bins used in the health analysis. In additional specifications we test smaller distance bins and continuous distance. We cluster standard errors at the plant level.[7]

3.2 First Closings

Many counties are nearby multiple units that close during our study time period. Using each closing as a unique event would introduce bias into our estimation because many months would enter into our analysis multiple times, sometimes as pre-closing and sometimes as post-closing. Instead, we choose to define the closing event for each county as the first closing within a 45 mile radius of the county's population centroid in the years 2005–2017.[8] This empirical strategy means we only capture the reduced form effect of a closing. Because closings are serially correlated, we capture both the effect of the original closing and the effect of subsequent closings. On average, each "first closing" represents the closing of 2.08 units within a plant because at many plants multiple units close in close proximity. Additionally, each county is exposed to an average 1.1 closing plants in the years following the first shutdown because closings among nearby plants are serially correlated. Thus, each closing should be thought of as the average effect of 2.29 units closing over the two years following first shutdown.

This allows for two interpretations of our results. First, they can simply be interpreted as the reduced-form effect of a county's first unit shutdown. Second, they can be scaled down by 2.29 to find the average effect of any unit shutdown under the assumption that later shutdowns have the same effect as earlier ones. For clarity of language, we refer to the effect of a first unit closing over the subsequent sections, but it is important to remember that this actually represents 2.3 units closing over the time period we are studying.

We consider separately plants that had all units close and partial closings of plants that had a unit close while others remained in operation. A plant is defined as partially closed if at least one generator is still operating in the year following the retirement year. We identify continued operation by looking at units that continue to report operating capacity and air pollution. The average plant had two operating units in 2004.

Unit of Analysis

In the health analysis, we use counties as the primary unit of analysis as counties are the most granular geographic level for which we have mortality data. For each county, we identify the first coal-fired generator closure within a 45 mile radius of the county's 2000 population centroid. Our choice follows Clay, Lewis, and Severnini (2016) who study the opening of US coal plants in the postwar era. Clay, Lewis, and Severnini (2016) assert that the majority of pollution will fall within 30 miles (50 kilometers) of a coal unit. Accordingly, we use counties 30–45 miles away from the closing unit as a control group.

Because counties have large differences in population, we choose to population weight our health results using 2005 county population. We do this for three reasons—first, population-weighting provides the policy-relevant parameter as it identifies the effect of closure on the average affected person rather than the average affected county. Second, this person-level effect allows for a more direct comparison with our housing results. Third, population-weighting improves our ability to distinguish an underlying mortality effect from random month-to-month variation in mortality. Counties with low population have substantial month-to-month variation in mortality rates because mortality is a relatively rare event in small counties. Giving these counties the same weight as large counties would introduce substantial noise into our analysis.

In the housing analysis, we use a separate transaction-level dataset that allows us to greatly improve the geographic precision of our results. We are able to identify home sales at the coordinate level and examine effects of closure both within distance bins and in continuous distance. The housing results give equal weight to each transaction. To the extent that transactions happen at a relatively constant frequency, this approximates population weighting.

4 Results

Generation and Pollution

We begin by examining the effect of a coal unit retirement on local generation and pollution. Although a shutdown by definition decreases generation and pollution at that unit, a county could see no net change (or even an increase) if other units nearby (or in the same plant) increased generation in response to the shutdown. In particular, we are interested in differences in exposure to generation by a county's distance to the retiring unit.

Accordingly, for the purposes of this analysis, our outcome variables are generation and pollution originating from units within 30 miles of a county population centroid, where the literature suggests much of the pollution harms takes place (Burney, 2020; EPA, 2005). We then compare changes in these outcome variables before and after closing between counties that are 0–30 miles away from the closing unit, and counties that are 30–45 miles away from the closing unit (omitted value).

Table 1 shows the results of estimating equation 3.1 with generation as an outcome. When we examine changes in coal generation within 30 miles of a unit closure, we observe declines in exposure to coal generation of roughly 100,000 MwH/month relative to counties 30–45 miles from the plant, or about 25% of average coal generation. However, this decline is entirely compensated by increased natural gas production; total power generation within 30 miles of the plants see little meaningful change.[9]

This fact is important for interpretation; rather than the effect of removing a generating unit with all else held constant, all estimates should be thought of as the reduced-form effect of a shutdown, which includes any compensating generation from surrounding natural gas plants to make up for the loss in capacity. This is also the policy relevant parameter. As these results show, shutdowns do not happen in a vacuum, but instead cause other local units to increase output. The policy-relevant mortality and housing effects are not simply the effect of the shutdown, but the effect of the shutdown net the effect of new nearby generation sources that come online to make up for the lost capacity. This effect is precisely what we are estimating here.

In Table 1, Columns 4 and 5 show the change in pollution caused by the shutdown. Despite the total level of generation within 30 miles of a closing unit remaining relatively similar, SO_2 and NO_X emissions from electricity generation units within 30 miles of the population centroid fall by around 60%–70% after closure relative to control counties. This reflects the fact that much of the compensating generation is from natural gas, which has much lower emission rates.

Appendix Figure A.3 shows these results in event study form. There are two important takeaways from these figures. First, there does appear to be some ramp-down in the months prior to shutdown. Beginning about a year before unit closure, we see decreases in generation and pollution. Second, the size of this ramp down is dwarfed by the large decline that occurs around the closure date. Importantly, any ramp down would lead to an underestimation of mortality effects as the difference between pre and post closure is smaller than the true difference in pollution exposure over a longer time horizon.

Adult Mortality

We next examine the effect of these generation and pollution declines on adult mortality. Table 2 shows our primary results. The even columns show the effects on the mortality rate, while the odd columns show the effect on the log transformation of the mortality rate. Both specifications show that a unit closing significantly decreases cardiovascular mortality. Being within 15 miles of a closing unit leads to a .46 cardiovascular deaths/100,000 population reduction in mortality (a 2.8% decline, $p < .01$), while being within 15–30 miles of a closing unit leads to a .35 fewer cardiovascular deaths/100,000 population (a 1.5% decline, $p < .05$). There appears to be little effect on overall respiratory mortality; this may be because many respiratory deaths are reflective of longer-term exposure. Respiratory deaths are also a rarer event, making a change harder to detect.

Columns (7) and (8) show the effect of the coal unit closure on non-respiratory, non-cardiovascular mortality. Pollution from coal plants largely affects respiratory and cardiovascular health in the short run, thus we would not expect closure to have any effect on non-cardiovascular, non-respiratory deaths. That is precisely what we see here; the effects of closure are small, statistically insignificant and do not vary systematically with distance. This provides additional support that the observed effects are due to unit closures and not other factors correlated with both closures and overall death rates.

Figure 1 shows the effects of closure on cardiovascular, respiratory, and all other mortality in event study form for counties less than 30 miles (relative to being 30–45 miles from the closing unit) before and after unit shutdown. Prior to unit shutdown, the effect on cardiovascular mortality is consistently small, insignificant and demonstrates no clear pre-trend. After closure, we can see an immediate downward shift in mortality, which persists for much of the rest of the period. Although estimates are quite noisy, the shift in mortality at the month of closure and the lack of a pre-trend nevertheless provides additional evidence that the observed results are unlikely to be driven by differential underlying trends in mortality between close and far counties. The figures for respiratory and non-respiratory, non-cardiovascular mortality show no clear trends, again supporting the idea that counties exposed to unit closures did not have differential trends in mortality prior to shutdown.

[[Insert Figure 1 about here]]

Appendix Table A.2 tests the robustness of these effects to the inclusion of a large number of covariates. Columns (1) and (5) include plant-specific time trends. Columns (2) and (6) include 2004 income and population variables interacted with year-month to control for any differential mortality trends that may occur in counties with demographic or economic profiles. Columns (3) and (7) add baseline generation, sulfur dioxide pollution and NO_X pollution interacted with year X month to control for any differential trends that might occur between counties with different underlying exposure to generation. Columns (4) and (8) control for 2004 respiratory mortality and cardiovascular mortality interacted with year X month to control for different trends that might occur between counties with different underlying baseline mortality levels. The coefficients on both cardiovascular and respiratory mortality remain nearly the same in each of these four specifications, increasing confidence that our results are not driven by underlying differences between counties.

Appendix Table A.3 shows the robustness of estimates to different weighting schemes and bandwidths. Columns (1) and (5) show the results from unweighted regressions. Consistent with the expected measurement error in mortality rates for small counties, observed effects decline slightly and standard errors increase appreciably, but the effect on cardiovascular mortality remains highly economically (but not statistically) significant (0.38 death decline/100,000 population, p=0.131). Columns (2) and (6) show the results using a balanced panel; the results remain largely unchanged.[10] Columns (3) and (7) show the results using a 1 year bandwidth. The effects on cardiovascular mortality and respiratory mortality are roughly the same. Columns (4) and (8) show results using a 3 year bandwidth. Both cardiovascular and respiratory mortality effects shrink slightly, but remain economically significant. The decrease in cardiovascular mortality remains significant for counties within 15 miles of a closing unit.

Finally, Appendix Table A.4 shows these effects differentially by age, broken down by individuals over-65 and under-65. Cardiovascular mortality effects appear to be roughly uniform across age groups in percentage terms. Both over-65 and under-65 individuals within 15 miles of the closing unit see cardiovascular mortality fall by roughly 2%, although the effect is only statistically significant for elderly individuals.[11] Interestingly, unit closure does appear to lead to declines in respiratory mortality for middle-aged individuals, but not the elderly. Individuals under 65 living in a county within 15 miles from the closing unit see a 0.07/100,000 population decrease in respiratory mortality or a 4% decline.

Economic Changes

Unit closings may affect local economies in two principal ways that may affect housing values. First, coal unit closures may lead to unemployment among unit workers. However, since coal plants are relatively capital intensive, unit closures are unlikely to have a meaningful effect on employment in all but the most rural areas. Second, if a coal plant is linked to a nearby mine, unit closures may lead to negative demand shocks for local coal, leading to increased unemployment among miners.[12] Again, this is a largely rural phenomenon and, because we weight by population, is unlikely to affect our main results.

Table 3 shows the primary results of this analysis. Wages and employment effects are uniformly small, relatively precisely estimated and do not vary systematically with distance, suggesting that economic effects are not driving the observed results.[13] There is some suggestive evidence that the counties closest to the plant see a small ($\approx$ 4%) decline in utility and transportation employment, but this effect is both highly insignificant and much too small to have a meaningful effect on local economic health.

We next consider the possibility that coal plant closures may have had effects on nearby coal mines. Using quarterly data on mining employees and production by county, we test this hypothesis by comparing county-level mining outcomes before and after unit closure in counties close to the closure relative to those 30–45 miles away. As Table A.7 shows, there is no evidence that local mining production, local mining employees or local mining hours fell after unit closure. All point estimates are both economically and statistically insignificant.

In general, it appears that unit closures do not cause economic changes in a way that could explain the capitalization results that we present in the next section. These results also provide additional support for our empirical strategy since underlying differences across near and close counties (and counties experiencing differential retirement timings) do not differ in their economic experiences during the coal transition.

Home Prices

To examine the extent to which the coal unit retirement effects (i.e. health effects) are capitalized into home prices, we look at how home sales prices evolve following unit retirements. Figure 2 shows a monthly event study of home price changes comparable to the analyses in Appendix Figures A.3 and 1. For homes within 15 miles of a plant, we see prices start to increase 6–10 months following unit retirement.[14] Appendix Figure A.5 considers the event study over a longer time horizon. This figure plots the effect of log distance from plant interacted with year relative to unit retirement. Following closure, homes closer to plants appreciate relative to those farther away in a manner that persists up to 8 years after the unit retirement.

[[insert Figure 2 about here]]

Table 4 shows our core housing results based on the main specification described in Equation 3.3 using a repeat sales analysis. As can be seen in Column 1, the retirement of at least one coal unit at a plant leads to an approximately 3% increase in home prices for homes within 15 miles of the retiring plant. Columns 2 and 3 then illustrate that this effect is entirely driven by full retirements, which bump the housing appreciation figure to 5%. Despite the pollution reductions and health improvements[15] associated with partial closure, we do not see a home price effect for partial closure. Similarly, homes within 15–30 miles of the retiring unit, which experience considerable health improvements as a result of the closure, experience no commensurate change in housing values.[16] This pattern of results remains unchanged, with strikingly similar coefficients, when based on a cross-sectional comparison (see Appendix Table A.11).[17]

Further confirming that effects are concentrated very close to the retiring plant, Appendix Table A.12 repeats the analysis of Table A.11 using smaller bins. The effect of closure is concentrated among houses within 20 miles of the plant, and the effect for plants less than 10 miles away is statistically significantly larger than that for houses in the 10–20 mile bin.[18]

When interpreting these results, it is important to note that we have not accounted for any anticipatory effects of a shutdown. For example, if a shutdown was credibly announced in advance and potential homebuyers had access to this information we might expect the impact of the closure to already be partially reflected in home prices by the time the closure occurred. We believe such anticipatory effects are unlikely because accurate information about shutdowns is often in short supply.[19] However, even if such anticipatory effects exist, they would likely only bias our home price estimates towards zero since some of the price effect would have occurred prior to the closure; our estimates can be thought of as a lower bound.

Valuing Health

We are now ready to calculate the implicit value that households place on a statistical life. To be clear, this exercise requires a number of untestable assumptions and should be viewed as illustrative, with interpretation of point estimates treated with a great deal of caution. We begin with the numbers for cardiovascular and respiratory deaths based on column 5 of Table 2, which indicate that coal unit shutdowns leads to 0.51 fewer cardiovascular and respiratory deaths/100,000 population per year among counties 0–15 mi from the shutdown and 0.315 fewer cardiovascular deaths/100,000 population per year among counties 15–30 mi from the shutdown.[20] Since this risk will be priced at the household level, we need to convert this individual annual risk into a household one. The average household size in the U.S is 3.4 (Fry, 2019) so this translates into a total reduction in mortality risk of 0.00001734 and 0.00001071 per household per year, respectively.

Our calculation of the corresponding increase in housing values is based on the results summarized in Table 4. Here we see that average housing values are $217,535 for homes with 15 miles of retiring unit and $233,046 for homes between 15 and 30 miles of the closing unit. For those close to the plant, housing values increase by approximately 3 percent while those further away appreciate at as statistically insignificant 0.3 percent. Combining these figures suggest that homes with 15 miles of the decommissioned coal unit appreciated by $6526 while those further away experience an imprecisely estimated increase of $649. Since the health risks are annual, we we need to convert this capitalization value into an annual cost. We do so by annuitizing that value into a price based on the opportunity cost of capital of 5% (the average mortgage rate during our study period). In the end, this translates into an annual price per household of $326 for those within 15 miles of the plant and $32 for those between 15 and 30 miles of the plant.

Armed with these annual risk reductions and additional costs per household, we can provide a rough estimate for the willingness to pay to avert death, or the value of a statistical life. Our results suggest that those homeowners within 15 miles of the plant have a VSL of $18.8M. In contrast, those with homes between 15 and 30 miles of the plant have an imprecisely estimated VSL of $0, or if we utilize the statistically insignificant point estimate of housing price effects for this group then a VSL of $2.99M.

This estimate for those within 15 miles of a plant is approximately double the estimate used by EPA (EPA, 2015) while the figure for those within 15–30 miles is quite close to the lower end of the range found in a meta-analysis summarizing a range of estimate in the literature (Viscusi and Aldy, 2003). More importantly, this result really drives home the important role of perceptions when using revealed preference data to determine the value of statistical life. While the improvement in health risk for those further from the plant is slightly smaller than for those immediately adjacent, the big driver of these differences in how those risks capitalize. Our findings are consistent with the notion that those further from the plant are largely unaware of the health benefits and thus do not price them into housing values.

5 Conclusion

The past two decades have witnessed an unprecedented transition from coal to gas-fired powerplants in the United States. We exploit data on this transition to examine the health and economic impacts of these changes and estimate how those changes capitalize into housing values. Using a difference-in-differences design, we find that individuals in counties whose population centroid is within 30 miles of a plant that closes at least one coal-fired unit experience 1.5–2% declines in cardiovascular mortality following shutdown. While these health improvements appear to capitalize into housing values, they only do so for homes within 15 miles of the plant. Based on our back-of-the-envelope calculation, our results suggest that those within 15 miles of the plant have a value of statistical life (VSL) of $18.8M while those between 15–30 miles of the plant have a VSL of $2.99M. Given the absence of employment or wage effects from shutdowns, the most plausible explanation for these differences in capitalization is differences in perceived risks. This point is reinforced by the fact that capitalization effects only occur when coal-unit retirements are complete rather than partial, suggesting that the visibility of these changes may play an important role in pricing outcomes.

This idea of inattention and risk salience is well established in the psychological and behavioral economic literature (e.g. Johnson and Tversky, 1983; Slovic, 1995; Bordalo, Gennaioli, and Shleifer, 2012; Gabaix, 2019), but appears under-explored in the environmental economics and hedonics literature. Moreover, this pattern of results does not appear to be unique to our setting. Recent work on air toxins using a very different empirical strategy has also found that health impacts extend beyond the geographic range of capitalization around the locations of toxic plants (Currie et al., 2015), though this geographic mismatch was not the focus of their paper.

On the one hand, these finding are quite intuitive. Buyers of homes many miles from a pollution source are unlikely to view them as a factor when bidding on a property. Indeed, we implore the reader to name the major pollution sources more than 15 miles from their current residence. On the other hand, it is important to note that it is not necessary to know the source of pollution to have it priced into a property. Knowledge of the levels of pollution in a given location should be sufficient for capitalization. Yet, despite the considerable availability of publicly available data on air quality and the proliferation of cheap residential air quality monitors that can provide hyper-localized information, this does not appear to have reached a large enough audience to make a significant impact on housing prices for those properties further away from the primary pollution sources.

This raises a number of important questions. First, if perceived risk plays an important role in shaping housing prices across geographies, along what other dimensions might this vary? How might education or other demographics alter the wedge between subjective and objective beliefs? Second, why is this heterogeneity not resolved through market arbitrage, such that we settle on a singular risk premium that lies somewhere between belief extremes? Finally, if hedonic estimates are highly sensitive to differences in inter-individual ‘experiences’, then what does this imply about cross-sectional approaches to estimation? Might they help explain some of the inconsistencies in the literature (Smith and Huang, 1995)? Together, these questions comprise a future research agenda. Regardless, it is clear that much more work is needed to better understand the perception production function and how that influences the standard interpretation of housing hedonic studies as a measure of willingness to pay for non-market goods (Rosen, 1974). It also points to potentially fruitful collaborations between environmental and behavioral economists.

Acknowledgements

We thank UCSD seminar participants for many useful comments. We are grateful to Jen Burney for sharing her knowledge of pollution transmission and for making this a better paper. Sam Krumholz acknowledges (and thanks) the Sloan Foundation for support during the writing of this article. Rebecca Fraenkel acknowledges (and) thanks the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1144086. We are also grateful to Zillow for providing access to their housing data through the Zillow Transaction and Assessment Dataset (ZTRAX) and ZTRAX workshop participants for their input on how to best use this data. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, Zillow, the Department of Justice or the National Bureau of Economic Research.

References

[1] Michelle L Bell et al. “Ozone and short-term mortality in 95 us urban communities, 1987-2000.” In: *Jama* 292.19 (2004), pp. 2372–2378.

[2] Kelly C Bishop et al. “Best practices for Using Hedonic Property Value Models to Measure Willingness to Pay for Environmental Quality”. In: *Review of Environmental Economics and Policy* (2020).

[3] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. “Salience Theory of Choice Under Risk”. In: *The Quarterly Journal of Economics* 127.3 (2012), pp. 1243–1285.

[4] Eli Brewer et al. “PM2. 5 and Ultrafine Particulate Matter Emissions from Natural Gas-fired Turbine for Power Generation”. In: *Atmospheric Environment* 131 (2016), pp. 141–149.

[5] Jennifer A Burney. “The Downstream Air Pollution Impacts of the Transition from Coal to Natural Gas in the United States”. In: *Nature Sustainability* 3.2 (2020), pp. 152–160.

[6] Brantly Callaway and Pedro HC Sant’Anna. “Difference-in-differences with multiple time periods”. In: *Journal of Econometrics* 225.2 (2021), pp. 200–230.

[7] Rahil Changotra et al. “Comparative Assessment and Validation of USEPA Air Quality Modelling Techniques Using Pre-Post Point-Source Observations Near a Coal-Based Thermal Power Plant”. In: *Bow Ties in Process Safety and Environmental Management* (2021), pp. 181–198.

[8] Kenneth Y Chay and Michael Greenstone. “Does Air Quality Matter? Evidence from the Housing Market”. In: *Journal of Political Economy* 113.2 (2005), pp. 376–424.

[9] Karen Clay, Joshua Lewis, and Edson Severnini. *Canary in a Coal Mine: Infant Mortality, Property Values, and Tradeoffs Associated with Mid-20th Century Air Pollution.* Tech. rep. National Bureau of Economic Research, 2016.

[10] Congressional Research Service. "Cost and Benefit Considerations in Clean Air Act Regulations". In: *Report #R44840* (2017).

[11] Janet Currie et al. "Environmental Health Risks and Housing Values: Evidence from 1,600 Toxic Plant Openings and Closings". In: *The American Economic Review* 105.2 (2015), pp. 678–709. issn: 00028282. URL: `http://www.jstor.org/stable/43495396`.

[12] Rebecca J Davis, J Scott Holladay, and Charles Sims. *Coal-fired power plant retirements in the US*. Tech. rep. National Bureau of Economic Research, 2021.

[13] Clément De Chaisemartin and Xavier d'Haultfoeuille. "Two-way fixed effects estimators with heterogeneous treatment effects". In: *American Economic Review* 110.9 (2020), pp. 2964–2996.

[14] Francesca Dominici et al. "Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases". In: *Jama* 295.10 (2006), pp. 1127–1134.

[15] EPA. *Health and Environmental Effects of Particulate Matter*.

https://www.epa.gov/pm-pollution/health-and-environmental-effects-particulate-matterpm. Accessed: 2018-03-30. 2018.

[16] US EPA. “Regulatory Impact Analysis of the Final Revisions to the National Ambient Air Quality Standards for Ground-level Ozone”. In: *EPA-452/R-15-007* (2015).

[17] US EPA. “Revision to the Guideline on Air Quality Models: Adoption of a Preferred General Purpose (Flat and Complex Terrain) Dispersion Model and Other Revisions; Final Rule”. In: *Federal Register* 70.216 (2005), pp. 68218–68261.

[18] Harrison Fell and Daniel T Kaffine. “The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions”. In: *American Economic Journal: Economic Policy* (2017).

[19] Rebecca Fraenkel and Samuel Krumholz. *Property Taxation as Compensation for Local Externalities: Evidence from Large Plants.* Tech. rep. Unpublished working paper, 2022.

[20] R. Fry. *The number of people in the average U.S. household is going up for the first time in over 160 years.* 2019. URL: \url{https://www.pewresearch.org/facttank/2019/10/01/the-number-of-people-in-the-average-u-s-household-is-going-up-for-the-first-time-in-over-160-years/}.

[21] Xavier Gabaix. “Behavioral Inattention”. In: *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 2. Elsevier, 2019, pp. 261–343.

[22] Andrew Goodman-Bacon. “Difference-in-differences with variation in treatment timing”. In: *Journal of Econometrics* 225.2 (2021), pp. 254–277.

[23] Josh Graff Zivin and G Singer. *Disparities in Pollution and Capitalization Rates*. Tech. rep. UCSD, 2022.

[24] Michael Greenstone and Justin Gallagher. “Does hazardous waste matter? Evidence from the housing market and the superfund program”. In: *The Quarterly Journal of Economics* 123.3 (2008), pp. 951–1003.

[25] Kevin Haninger, Lala Ma, and Christopher Timmins. “The Value of Brownfield Remediation”. In: *Journal of the Association of Environmental and Resource Economists* 4.1 (2017), pp. 197–241.

[26] Akshaya Jha and Nicholas Z Muller. “Handle With Care: The Local Environmental Costs of Coal Piles at US Power Plants”. In: (2017).

[27] Reid Johnsen, Jacob LaRiviere, and Hendrik Wolff. “Fracking, Coal and Air Quality”. In: (2017).

[28] Eric J Johnson and Amos Tversky. "Affect, Generalization, and the Perception of Risk." In: *Journal of Personality and Social Psychology* 45.1 (1983), p. 20.

[29] Klea Katsouyanni et al. "Short term effects of ambient sulphur dioxide and particulate matter on mortality in 12 European cities: results from time series data from the APHEA project". In: *Bmj* 314.7095 (1997), p. 1658.

[30] Sam Krumholz. "The Effectiveness of Litigation as a Policy Instrument: The Case of the New Source Review Settlements". In: *Available at SSRN 3398035* (2019).

[31] Joshua Linn and Kristen McCormack. *The Roles of Energy Markets and Environmental Regulation in Reducing Coal-Fired Plant Profits and Electricity Sector Emissions.* Resources for the Future, 2017.

[32] Darrell Proctor. "West Virginia Coal Plant Will Remain Open Until 2022". In: *Power* (Oct. 19, 2018). url: `https://www.powermag.com/west-virginia-coal-plant-will-remain-open-until-2022/` (visited on 04/24/2023).

[33] Sherwin Rosen. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition". In: *Journal of Political Economy* 82.1 (1974), pp. 34–55.

[34] Paul Slovic. "The Construction of Preference". In: *American Psychologist* 50.5 (1995), p.

364.

[35] V Kerry Smith and Ju-Chin Huang. “Can Markets Value Air Quality? A Meta-analysis of Hedonic Property Value Models”. In: *Journal of Political Economy* 103.1 (1995), pp. 209–227.

[36] W Kip Viscusi and Joseph E Aldy. “The Value of a Statistical Life: a Critical Review of Market Estimates Throughout the World”. In: *Journal of Risk and Uncertainty* 27.1 (2003), pp. 5–76.

Table 1: Effect of unit closure on generation and pollution

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Coal Gen	NG Gen	Total Gen	SO2 Tons	NOX Tons
Post X $<$ 30 mi	-107855.9^{***}	97955.0^{*}	-39321.4	-1093.3^{***}	-322.8^{***}
	(26355.6)	(57361.7)	(65488.2)	(333.2)	(103.7)
Observations	36576	36576	36576	36576	36576
R^2	0.954	0.909	0.937	0.890	0.900
Population Weighted	Y	Y	Y	Y	Y
Dep. Var. Mean	463928.1	260620.2	761301.9	1341.7	591.1

Clustered standard errors in parentheses

$^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

This table shows the effects of a coal unit closure within 30 miles of a county population centroid on coal generation (MwH/month), natural gas generation (MwH/month), total generation (MwH/month), SO_2 tons produced and NO_X tons produced within a 30 mile radius of the county population centroid in the two years before and after a coal unit closure. We include controls for county fixed effects and year X month X Census division fixed effects. All standard errors are clustered at the closing plant level and regressions are weighted by 2005 county population.

Table 2: Effect of unit closure on local mortality

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Tot Card	ln(Card Mort)	Resp Mort	ln(Resp Mort)	Card + Resp	ln(Card + Resp)	All Other Mort	ln(All Other)
Post X 15–30 mi	-0.353^{**}	-0.0147^{**}	0.0387	0.0188	-0.315^{**}	−0.00750	-0.000345	0.00496
	(0.142)	(0.00678)	(0.0783)	(0.0119)	(0.146)	(0.00584)	(0.221)	(0.00558)
Post X <15 mi	-0.463^{***}	-0.0277^{***}	−0.0471	−0.00622	-0.510^{**}	-0.0219^{***}	0.184	0.00650
	(0.158)	(0.00801)	(0.0786)	(0.0137)	(0.188)	(0.00760)	(0.311)	(0.00700)
Observations	36549	35950	36549	32898	36549	36272	36549	36420
Dep. Var. Mean	21.93	3.022	6.627	1.786	28.56	3.290	40.09	3.648
Population Weighted	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

$^{*}p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

This table shows the effects of a coal unit closure on total cardiovascular mortality rate, total respiratory mortality rate, total cardiovascular + respiratory mortality rate, and all-other-cause mortality rate within 45 mile radius of the county population centroid in the two years before and after a coal unit closure. Mortality rates are per 100,000 population. All distance bin indicators are relative to the 30–45 mile distance to closing unit category. We include controls for county fixed effects and year X month X Census division fixed effects. All standard errors are clustered at the state level and all regressions are population weighted. Respiratory deaths are defined as deaths with an ICD-10 code prefix of “J”, cardiovascular deaths are defined as deaths with an ICD-10 code prefix of “I.”

Table 3

Effect of unit closure on local economic outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
	Mean Empl.	Mn Empl.	Ln(Mine+Util Emp.)	Avg. Wage	Ln Avg. Wage	Ln(Mine +Util Wage)
Post X 15-30 mi	2371.8	0.000449	0.00416	3.717	0.00230	−0.00104
	(4857.4)	(0.00329)	(0.0709)	(2.493)	(0.00201)	(0.0171)
Post X <15 mi	3577.8	−0.00173	−0.0473	1.679	0.000582	−0.00612
	(4582.3)	(0.00297)	(0.0813)	(2.831)	(0.00249)	(0.0219)
Observations	13185	13183	6872	13185	13183	6872
Adjusted R^2	1.000	1.000	0.952	0.928	0.967	0.815
Dep. Var. Mean	391964.1	12.00	6.382	974.7	6.849	7.359
Population Weighted	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

$^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

This table shows the effects of a coal unit closure on economic outcomes within a 45 mile radius of the county population centroid in the two years before and after a coal plant closure. Ln employment and ln wage are measured quarterly and taken from the Quarterly Census of Employment and Wages (QCEW). Ln population and ln median income are measured annually and taken from the Census’ Small Area Income and Poverty Estimates (SAIPE). All distance bin indicators are relative to the 30–45 mile distance to closing unit category. We include controls for county fixed effects and year X month X Census division fixed effects. All standard errors are clustered at the county level and all regressions are population weighted.

Table 4

Home Price Effects by Full vs Partial Closure: Repeat Sales

	(1)	(2)	(3)
	Ln Price	Ln Price	Ln Price
<15 mi × Post Retire	0.0312 **	0.0516 ***	-0.0107
	(0.0142)	(0.0173)	(0.0190)
15–30 mi × Post Retire	0.00372	0.0128	-0.00725
	(0.0105)	(0.0141)	(0.0129)
Partial	Both	No	Yes
N	5257440	2991263	2266177
# of clusters	73	42	31
Adjusted R2	0.775	0.762	0.771
Mean Px 0-15	217535.5	233046.8	181954.5
SD 0-15	185535.2	199305.0	143002.9
Mean Px 15-30	216462.7	252215.1	172504.8
SD 15-30	174282.1	183642.1	150835.3
Mean Px 30-45	190994.7	220281.0	163023.0
SD 30-45	154611.4	175503.4	125363.2
N15	1342403	934857	407546
N30	2284839	1260022	1024817
N45	1630198	796384	833814

Standard errors in parentheses

$^{*}\ p < 0.10,\ ^{**}\ p < 0.05,\ ^{***}\ p < 0.01$

This table shows the effect of a coal unit retirement on home prices. Homes within .5–45 miles of the plant and transactions from 6 years before to 8 years after the retirement are included. Regression includes year X month X plant fixed effects. Standard errors are clustered at the plant level.

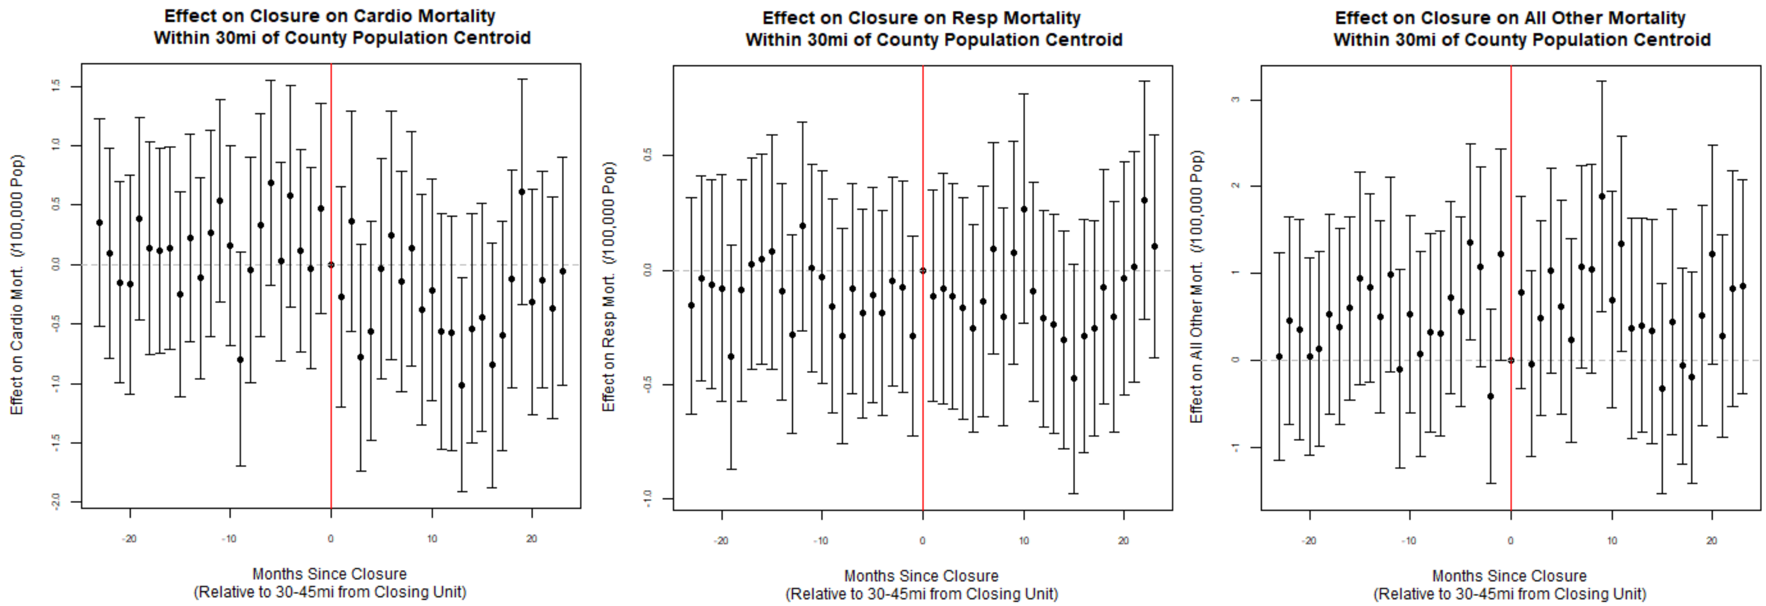

Figure 1

Changes in Mortality For Counties Experiencing Closure in the Two Years Before and After the Closure

Note: All coefficients come from a regression of the outcome variable on a months since closure indicator interacted with a distance bin indicator for less than 30 miles from closing plant (relative to 30–45 miles), a year X month X Census division fixed effect and county fixed effects. Respiratory deaths are defined as deaths with an ICD-10 code prefix of “J”, cardiovascular deaths are defined as deaths with an ICD-10 code prefix of “I.” All regressions are weighted by 2005 county population and all standard errors are clustered at the county level.

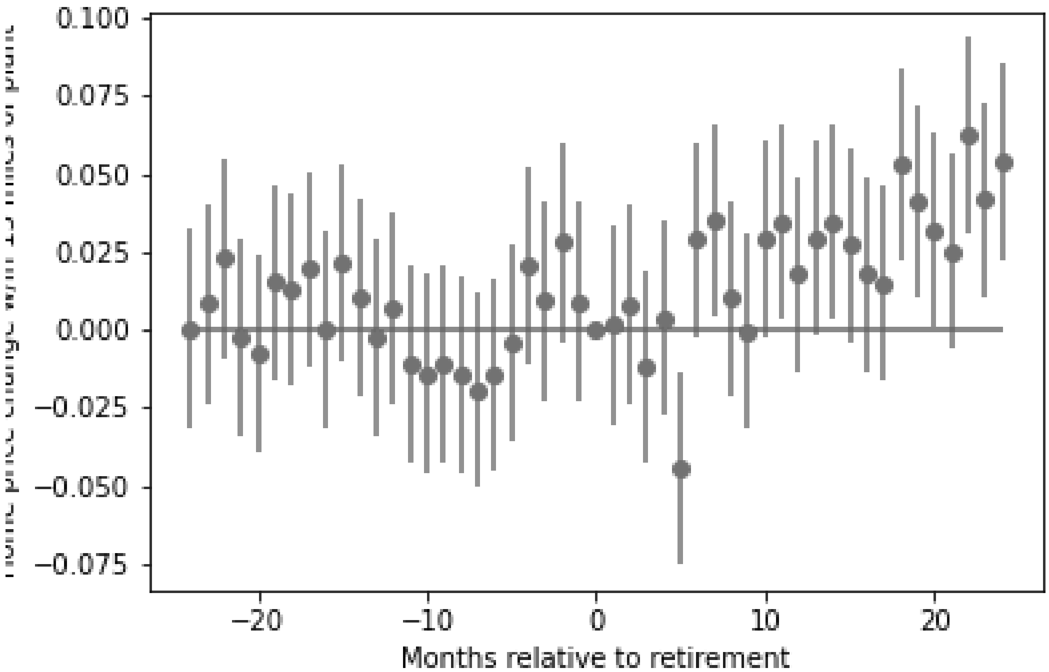

Figure 2

Home Price Effects of Unit Retirement: 0–15 miles vs 15–60 mi

Note: This shows how home prices change for repeat sales within 2 years of a closure. All coefficients come from a regression of log home price on the interaction of being within 15 miles of a plant and indicators for months since closure and property fixed effects. Homes .5–60 miles from a retiring plant are included.

NOTES

1. The average plant had 2.1 units in 2004.
2. The average plant had 2.1 units in 2004.
3. We use county of residence and not county of death as the locator variable because we are interested in changes of death due to everyday exposure. County of death may be dictated by the hospital in which the deceased sought care.

4. We include home sales transactions in AL, AZ, CA, CO, CT, DE, FL, GA, IL, IN, IA, KY, MD, MA, MI, MN, NE, NV, NH, NJ, NY, NC, OH, OK, PA, RI, SC SD, TN, VA, WV, and WI. Appendix Figure A.4 shows plants that are included in the housing analysis.
5. This also restricts our sample to counties whose population centroid is within 45 miles of a coal unit closing during our sample period.
6. We use a balanced panel where counties in the 30–45 mile bin can be viewed as "never treated" with 2 years of pre and post data for each plant closure and present event study results to address concerns about weighting in DiD estimates driven by length of exposure to treatment raised by (Goodman-Bacon, 2021; De Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021).
7. We take several steps to address concerns about weighting in DiD raised by Goodman-Bacon, 2021; De Chaisemartin and d'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021: We limit our sample for each plant to home sales 6 years before and eight years after a plant closure to limit the influence of "early treated" units. Homes in the 30–45 mile bin can be viewed as "never treated" and time trends are plant specific. We present event study estimates analogous to those for the health results in Figure 2 and with no sample restrictions in Appendix Figure A.5.
8. Almost no coal units closed between 1998–2005.
9. This effect is driven by a relatively small number of plants. Around one-third of counties experience changes in natural gas generation after unit shutdown, but these changes are large.
10. This restricts our sample to plants that closed between 2007 and 2013.
11. This effect is similar for both groups in percentage terms but only significant for the elderly, possibly due to the higher baseline cardiovascular mortality.
12. For plants that convert to NG, there could also be a temporary *increase* in employment in construction.
13. Appendix Table A.8 repeats the analysis of Table 3 with 8 years of pre and post closure data included to check for longer term employment trends that could influence our housing analysis. We again see no change in wages and employment.
14. Theory would suggest that if unit closure acts as an information shock to prospective homebuyers, the home price effect of the closure should manifest itself immediately and then plateau. However, that is not the pattern observed here, instead we see that prices of homes near the closing plant increase steadily over the first 24 months following plant closure. This may

reflect a number of factors, including: the slow diffusion of information among home buyers; ongoing, but diminishing dis-amenities from plant closure (e.g. construction, transport of toxic materials); and subsequent unit closures following the first unit closure in a county. While we lack the data and statistical power to identify the mechanisms driving the home price dynamics we observe, we believe doing so is an important direction for future research.

15. See Appendix Table A.5.

16. We also do not see evidence of awareness of pollution intensity by home buyers. In table A.6, we see suggestive evidence that counties near plants with above median baseline (2004) NOx pollution have see greater health improvements following unit closure. In Appendix Table A.10, while only plants with above median baseline NOx pollution have statistically significant health impacts, the point estimate for below median plants is larger than it is for the more polluting plants.

17. Appendix Figure A.6 shows effects by plant. While the estimates are noisy, the trend of larger positive home price improvements for full closures is clear.

18. Appendix Table A.9 shows results broken out by whether the plant is in an urban or rural location. The results are suggestive of a smaller effect for urban plants but cannot be statistically significantly distinguished from one another due to the small number of urban plants.

19. One anecdotal example of this can be seen in a coal plant in West Virginia, which announced it would close in 2019, but then in October 2018 pushed the closing date out until 2022 (Proctor, 2018). Our understanding is that situations similar to this one are not uncommon and illustrate why an alternate specification of measuring the impact beginning with an “announced” retirement date would be challenging.

20. It is important to note that we are only attributing lives saved to a coal plant closure based on data from the two years following the closure.

21. Many transaction records only provide a month and year of sale. The 93 day window allows for any three month window regardless of month length.

22. Many events have multiple transactions recorded in the ZTRAX database due to mortgage changes, adjustments, multiple foreclosure notices, etc.

23. Our largest bins are 6+ bedrooms, 5+ baths. Square footage is in 500 square foot bins and lot size is in 1 acre bins. Prior to 2000, year built is binned by decade.