

East Side Story: Historical Pollution and Persistent Neighborhood Sorting*

Stephan Heblich

Alex Trew

Yanos Zylberberg

October 2, 2019

Abstract

Why are the east sides of formerly industrial cities often the more deprived? Using individual-level census data together with newly created historical pollution patterns derived from the locations of 5,000 industrial chimneys and an atmospheric model, we show that this results from the persistence of neighborhood sorting that first emerged during the Industrial Revolution when prevailing winds blew pollution eastwards. Past pollution explains up to 20% of the observed neighborhood segregation in 2011, even though coal pollution stopped in the 1970s. We develop a quantitative model to identify the role of neighborhood effects and relocation rigidities underlying this persistence.

Keywords: Neighborhood Sorting, Historical Pollution, Persistence.

JEL codes: R23, Q53, N00.

*Heblich: University of Toronto, CESifo, IfW Kiel, IZA, SERC; stephan.heblich@bristol.ac.uk; Trew: University of Glasgow; alex.trew@glasgow.ac.uk; Zylberberg: University of Bristol, CESifo, the Alan Turing Institute; yanos.zylberberg@bristol.ac.uk. We thank the editor and six referees for comments that greatly improved this paper, as well as Philipp Ager, Spencer Banzhaf, Nate Baum-Snow, Kristian Behrens, Davide Cantoni, Peter Egger, Klaus Desmet, Christian Dippel, Gilles Duranton, Oliver Falck, Ed Glaeser, Walker Hanlon, Hans Hvide, Hans Koster, Robert McMillan, Jos van Ommeren, Giacomo Ponzetto, Diego Puga, Joshua Lewis, Jeffrey Lin, Petra Moser, Steve Redding, Jorge De la Roca, Stuart Rosenthal, Max Satchell, Hannes Schwandt, William Strange, Leigh Shaw-Taylor, Jon Temple, Chris Timmins, Nico Voigtländer, and Ludgar Wössmann for useful discussions and comments. We also thank seminar participants in Bayreuth, Bergen, Bern, Boulder, Bristol, Cambridge, CERGE-EI, CUHK, Dublin, Durham, Edinburgh, Essex, HSE St. Petersburg, Kiel, Munich, NYU Stern, Odense, Passau, the Philadelphia FED, SERC, St Andrews, Stirling, Toronto, UCLA Anderson, USD, Warwick and participants at the 2016 IEB Workshop on Urban Economics, the 2016 EUEA, the 2016 UEA, the 2017 ASSA meeting and the 2017 NBER EEE meeting for helpful comments. We are grateful to a number of excellent research assistants. This work was part-funded by the ESRC through the Applied Quantitative Methods Network: Phase II, grant number ES/K006460/1. The usual disclaimer applies.

Cities that were formerly reliant on industry tend to have eastern suburbs that are notably poorer than western suburbs. This observation is echoed in media stories about the east side of London, New York or Paris and in popular culture (such as in the long-running BBC soap opera, *EastEnders*). We show that the east-west gradient is a remnant of the distribution of the atmospheric pollution which affected cities during the Industrial Revolution. Pollution from historical factories accounts for 15% of the variation in neighborhood composition in 1881. While industrial coal pollution effectively stopped in the 1970s, its repercussions are still felt today as shown by the dynamics of neighborhood segregation during the post-pollution period. We use a quantitative model of neighborhood sorting to show how the combination of neighborhood effects and relocation frictions generated tipping-like dynamics which anchored formerly polluted neighborhoods.

Although the impact of pollution on welfare in cities is often highlighted in policy debates, the long-run responses of economic agents are less well known. Providing such evidence is challenging since systematic air pollution monitoring on a fine spatial scale only started after industrial coal pollution markedly slowed down. To fill this gap in our knowledge, we develop a novel method of modelling historical pollution within English cities at the end of the 19th century. This allows us to analyze the effect of pollution on neighborhood dynamics during industrialization as well as its effects in the long run, even after the success of the 1968 Clean Air Act. As coal pollution created a legacy in the form of neighborhood segregation, a better understanding of the underlying dynamics holds important implications for the design of environmental policies in newly-industrialized economies. The analysis of the dynamics of persistence is also instructive about urban policies aimed at reducing spatial inequalities in post-industrial economies. To this end, we develop a dynamic model of neighborhood sorting. The model provides evidence for non-linear forces in the dynamics of segregation, as in [Card et al. \(2008\)](#), arising from the interaction of relocation rigidities with neighborhood effects.

Our paper is the first to present a long-run analysis of the effects of pollution on the *internal* structure of cities. Our empirical analysis combines a novel measure of pollution exposure within British cities during the time of the Industrial Revolution with unique panel data at the neighborhood level spanning nearly 200 years. Three methodological innovations permit us to generate these data. To model pollution exposure, we first develop a method to geolocate industrial chimneys from historical Ordnance Survey (OS) maps of the 70 largest metropolitan areas in England over the period 1880–1900. Second, we construct pollution maps by using the locations of chimneys in an atmospheric dispersion model (ADMS 5) that incorporates additional

information on terrain, prevailing winds and chimney characteristics (chimney dimensions, exit velocity, coal burning temperature). Third, we develop a new method to assign individual entries in the 1881 census to low-level administrative units (for our purpose, the 2001 Lower Super Output Areas, LSOA). This procedure allows us to study residential sorting in the nineteenth and twentieth centuries, within 70 metropolitan areas and at the level of 5,500 LSOAs.

Our findings show a strong effect of air pollution on the share of low-skilled workers in 1881. A pollution differential equivalent to the one between the 10% and 90% most polluted neighborhoods of Manchester would be associated with a gradient of 16 percentage points in the share of low-skilled workers.

An obvious identification concern is that industrial chimneys are not randomly allocated within cities. We address this issue in the following two ways. First, we condition our analysis on the distance to industrial chimneys and analyze neighborhood composition in all different directions relative to the chimney. This spatial differencing exercise reveals excess deprivation in 1881 along a narrow corridor downwind of industrial chimneys. Importantly, there is no excess deprivation in neighborhoods downwind from industrial chimneys before the rise of industrial coal use. Second, we present an instrumental variable strategy to address the concern that chimneys may have been selectively located upwind of poor areas. Specifically, we instrument the pollution pattern induced by actual chimneys with a predicted pollution pattern that exploits exogenous locations of pollution sources. The choice of exogenous pollution sources rests on the fact that steam engines need water for cooling (Maw et al., 2012). We exploit waterways in 1827 as exogenous location factor and use points located uniformly along these waterways to predict the actual pollution pattern. As we condition on distance to the waterways and exclude neighborhoods bordering the waterside, we exploit the difference between upwind and downwind neighborhoods at the same distance from potential factories located along waterways. The IV specification delivers similar qualitative results as the baseline OLS specification.

Having established that pollution caused neighborhood sorting in the past, we focus on the dynamics of persistence between 1971 and 2011. After the second Clean Air Act of 1968, pollution from coal burning abruptly decreased.¹ Nevertheless, the social composition of neighborhoods in 1971, 1981, 1991, 2001 or 2011 points to a substantial persistence in the effect of historical pollution on within-city distribution of low-skilled workers. Past pollution explains up to 20 percent of the observed

¹The first Clean Air Act was enacted in 1956 as a reaction to the Great Smog of 1952 in London. However, the second Clean Air Act in 1968 caused a much more pronounced drop in coal consumption.

neighborhood segregation whether captured by the share of low-skilled workers or house prices. The dynamics of persistence between 1971 and 2011 show patterns of non-linearities, with mean-reversion for intermediate values of within-city pollution and inertia for neighborhoods with extreme values of within-city pollution.

In order to quantify these non-linear dynamics, we develop a model of residential sorting with two types of households—low-skilled and high-skilled—differing only in their income. The model relies on two important assumptions. First, the possibility for households to move is staggered; each household may be able to relocate free of costs with an exogenous probability in each period. Second, there is an idiosyncratic preference shock, which is distributed along an extreme value distribution (as in [Bayer et al., 2016](#), for instance). The model delivers a simple dynamic equation which characterizes the relative demand for neighborhoods. This equation—the dynamic equivalent of hedonic pricing—can be estimated from the observation of aggregate objects such as the share of low-skilled households in each neighborhood between 1971 and 2011. We estimate this relative demand for neighborhoods using Generalized Method of Moments in which historical pollution is an instrument for current neighborhood composition and its subsequent evolution. We find that the persistence of sorting is not only tied to relocation frictions, but to its interaction with strong preferences for neighborhood composition. High inherited deprivation has a direct, backward-looking effect on the persistence of sorting, as captured by the estimated share of residents that are given the opportunity to move in each period. An additional effect arises from the (forward-looking) expectations of movers: they anticipate a slow adjustment of neighborhood composition, which markedly affects their current valuation of the neighborhood. This last effect drives the non-linear persistence of segregation.

Our paper contributes to different strands of the literature. First, our work is related to [Lee and Lin \(2018\)](#) who look at exogenous natural amenities as driver of neighborhood sorting. Our paper, in contrast, studies the consequences of a temporary disamenity. To the best of our knowledge, we are the first to show that the large, temporary pollution from industrial coal use modified the spatial organization of cities in the long run.² Related papers that look at pollution-induced sorting today include [Banzhaf and Walsh \(2008\)](#) and [Chay and Greenstone \(2005\)](#).³

²Two recent papers find similar patterns of persistence for (i) historical marshes in New York City ([Villarreal, 2014](#)) and (ii) historical street car lines in Los Angeles County ([Brooks and Lutz, 2019](#)). A large literature discusses the path dependence in economic activity across cities (see, among others, [Davis and Weinstein, 2002](#); [Bleakley and Lin, 2012](#)).

³There is also a broad body of literature on pollution exposure and its effect on productivity ([Graff Zivin and Neidell, 2012](#)), cognitive performance ([Lavy et al., 2014](#)), violent crime ([Herrnstadt et al., 2019](#); [Heyes and Saberian, 2015](#)), and health ([Graff Zivin and Neidell, 2013](#); [Anderson, 2019](#);

[Kuminoff et al. \(2013\)](#) provide a broader review of the residential sorting literature. Our argument further relates to [Depro et al. \(2015\)](#) who argue that neighborhood sorting, rather than environmental injustice, is the reason why poor households are more exposed to environmental disamenities. Finally, our paper shares a common theme with [Hanlon \(2019\)](#), who argues that coal-based pollution was a significant disamenity with a strong negative impact on city size in England during the industrialization, and with [Chen et al. \(2017\)](#) and [Freeman et al. \(2019\)](#) who document a similar correlation between pollution and residential choices *between* cities in China.

Second, we relate to a burgeoning literature on the gentrification of historic centers in U.S. cities ([Brueckner and Rosenthal, 2009](#); [Guerrieri et al., 2013](#); [Baum-Snow and Hartley, 2016](#); [Couture and Handbury, 2019](#)) which builds on previous research on the dynamics of segregation and tipping points ([Schelling, 1971](#); [Card et al., 2008](#); [Logan and Parman, 2017](#)). After the sharp decline in industrial pollution in English cities, formerly polluted neighborhoods remain the poor parts of town. Our quantitative analysis points to non-linearities in the dynamics of segregation: highly-polluted neighborhoods repel high-income residents even after pollution has waned, through neighborhood effects. We differ from the tipping-point literature in two dimensions. In our context, we mostly identify a social component behind segregation (in contrast to the literature on the United States that mostly focuses on ethnic considerations). Moreover, we exploit a temporary disamenity to explain the initial spatial distribution of residents instead of permanent differences across neighborhoods.

Third, we make several methodological contributions to quantitative research in economic history. Our first methodological contribution is to provide an algorithm that geolocates census entries in 1881 and could be applied to any historical census in most developed countries. The algorithm exploits the clustering among census entries to infer the geolocation of all residents from a small share of well-matched neighbors. Our second contribution is to provide a methodology to digitize historical maps and fully exploit them as extremely valuable sources of information. Related to this approach is work by [Hornbeck and Keniston \(2017\)](#) and [Siodla \(2015\)](#), who use historical maps to understand the effects of the great fires in Boston and San Francisco, and [Redding and Sturm \(2016\)](#), who use maps to document Second World War destruction in London. Our third contribution is to show the ability of state-of-the-art pollution models to estimate historical pollution.

The remainder of the paper is organized as follows. Section [1](#) provides elements

[Deryugina et al., 2019](#)) which relates to our research. Closely related historical assessments of the effect of coal use on health include ([Barreca et al., 2014](#); [Clay et al., 2016](#); [Beach and Hanlon, 2018](#)).

of context and describes data sources. The reduced-form evidence on neighborhood sorting and its persistence is discussed in Section 2. Section 3 develops a dynamic model of residential choice. The identification and estimation of the demand for neighborhoods are discussed in Section 4. Finally, Section 5 concludes.

1 Historical background and data

This section describes the historical setting, our main data sources, and the construction and validation of atmospheric pollution between 1880–1900 and neighborhood composition in 1817, 1881 and 1971–2011. We also provide suggestive evidence about the role of wind directions in generating spatial inequalities within cities.

1.1 Historical background

The start of the Classical Industrial Revolution is dated to around 1760 by the arrival of new technologies in key growth sectors such as textiles, iron and steam. However, important consequences of that revolution were not realized until much later and per capita growth rates did not accelerate until after 1830 (Crafts and Harley, 1992). Economic growth was accompanied by an energy transition, with coal emerging as the dominant energy source around 1840.⁴ The rapid energy transition from then on is illustrated in the left panel of Figure 1: there is a sharp acceleration of coal consumption between 1850 and 1910 and a stabilization until 1960. The Clean Air Acts of 1956 and 1968 led to another energy transition: these Acts introduced regulations that penalized, among other things, the emissions of grit, dust and “dark smoke” in cities and placed minimum height restrictions on chimneys. The Clean Air Acts led industry to shift away from coal to the use of cleaner energy sources such as oil, gas and electricity generated by power stations outside of cities. As apparent in Figure 1, these regulations, and the Clean Air Act of 1968, had an immediate and marked impact on coal consumption.⁵

The heavy reliance on coal between 1850 and 1960 generated unprecedented concentrations of sulphur dioxide, which scarred cities and their surroundings.⁶ Mosley (2013), for instance, conjectured a relationship between historical pollution and

⁴As Musson (1976) shows, power derived from water wheels remained important to early nineteenth century industry—steam power was not prevalent outside of textiles until after the 1870s.

⁵The early twentieth century saw a consolidation of industry with employment peaking at 46% in 1950 (Crafts, 2014). The decline in coal consumption preceded the massive de-industrialization, which occurred most rapidly in the 1980s when state-owned industries were privatized.

⁶We do not attempt to discuss the growing body of literature on pollution and individual health. Instead, we refer to a survey article by Graff Zivin and Neidell (2013) and specifically to historical assessments of the effect of coal use on health discussed in Barreca et al. (2014), Clay et al. (2016), and Beach and Hanlon (2018).

neighborhood sorting: “In Manchester, prevailing and strongest winds [blow] from the south west. This meant that when the dense sulphurous smoke left Manchester’s tall chimneys it usually moved north east, and this was to have a marked effect on the shaping of the city. [...] The poorest city dwellers were forced to live amongst the mills and factories in north-easterly districts [...] the better-paid among Manchester’s working classes might at least escape the worst of the smoke.” The negative impact of atmospheric pollution on cities and their hinterlands is also captured in a well-known case of micro-evolutionary change. The dominant form of the peppered moth (*Biston betularia*) at the start of the nineteenth century was the lighter form (*insularia*) as it was camouflaged against predation when on light trees and lichens. The first sightings of the darker form of the moth (*carbonaria*) in the industrial north of England were not until after 1848 (Cook, 2003). As pollution caused trees to blacken under layers of soot, the *carbonaria* emerged as the dominant form by the end of the nineteenth century. The decline in air pollution after the Clean Air Acts induced a rapid recovery of the *Biston betularia insularia* after 1970 (Cook, 2003).

Along with the structural transformation of the economy, the end of the eighteenth century saw a rapid rate of urbanization with workers from the countryside flocking into the emerging industrial cities (Shaw-Taylor and Wrigley, 2014). As shown in the right panel of Figure 1, the growth of cities started to decline after 1830 and steadily slowed down as the nineteenth century proceeded. By the end of the nineteenth century, the large cross-country migratory flows that marked the early Industrial Revolution had moderated significantly.⁷

The great movement into cities came with an increase in density that was overwhelming Victorian cities at the beginning of the 19th century. The characteristic “back-to-back” houses were put together with thin walls and no foundation or ventilation, located in walking distance to the new factories. The lack of suitable housing along with limited supply of clean water and sanitation created unhealthy urban slums plagued by diseases like cholera and typhoid with a notoriously low life expectancy (Clark, 1962). Only in the second half of the 19th century did the Public Health Acts of 1872 and 1875 begin to change the living conditions of the poor.

While poor working-class families were typically stuck in the inner residential areas (or urban slums) around the city center, middle and subsequently lower-middle classes families started separating their place of work and place of residence thus encouraging new housing development in suburbs at the fringe of the city. Wealthier suburbs were characterized by private residential gardens and spacious villas while

⁷Williamson (1990) and Ravenstein (1885) show that the portion of city growth due to migration declines over the nineteenth century and, by 1881, 75% of individuals in England and Wales resided in the county of their birth.

poorer suburbs were made up of long terraces of byelaw housing. As discussed in [Heblich et al. \(2018\)](#), the rise of the railways and the subway facilitated suburbanization with residential development spreading widely around London. However, this wide-sprawling development was specific to London. In other British cities, suburbs remained in walking distance from work places ([Kellett, 1969](#); [Lawton, 1972](#)).

The empirical analysis relies on the following observable characteristics of neighborhoods: (i) neighborhood composition in 1817, before the acceleration in coal consumption and around the decline in rural migration to urban centers; (ii) atmospheric pollution, neighborhood composition and urban structure around 1880–1900, slightly before the peak in coal consumption; and (iii) neighborhood characteristics between 1971 and 2011, after the abrupt decrease in atmospheric pollution.

1.2 Data sources and construction

This section provides a brief summary of the different data sources; a comprehensive description can be found in Appendix [C](#).

OS maps and geo-location of pollution sources Data on city structure and pollution sources are drawn from the Ordnance Survey (OS) maps—25 inch to the mile (1842–1952). These maps come at a 25 inch: 1 mile scale, by far the most detailed topographic mapping that covers all of England and Wales. The maps contain details on roads, railway, rivers, canals, public amenities, the outline of each building and their use.⁸ Most useful for our purposes, these maps clearly mark the locations of all factory chimneys, in a sign of the fastidiousness of Victorian mappers. The symbol for a chimney is either a small rectangle with an inner circle or a large white circle, drawn to scale. In most maps, a *Chy* is written to help identify these symbols. The variations in symbols prevent us from directly using a recognition algorithm (two examples of symbols are shown in Figure [2](#)). We select the wave of maps published between 1880 and 1900 and restrict the analysis to the 70 largest cities in England at the end of the nineteenth century, as evident from the 1907 Census of Production. These cities constitute a quasi-exhaustive snapshot of industry, and cover 60% of the total population in 1801 and 66% in 2011.

We go through all maps and mark slightly more than 5,000 chimneys with a recognizable symbol *X* and a unique numeric identifier. An example of the chimney-

⁸We use the OS maps to extract information about the city structure and its amenities, most notably its contour, the location of town halls, market halls, churches, schools, universities, parks, theaters, museums, churches and hospitals. We benefited from the excellent research assistance of Andreas Arbin, Nicholas Cheras, Tim Ciesla, Joshua Croghan, Qingli Fang, Joanna Kalemba, Aishwarya Kakatkar, Matthew Litherland, Filip Nemecek, Ondrej Ptacek and Sava Zgurov.

identification is provided in Figure 3. On this map fragment, two different chimneys can be identified. The red symbol X , located in the center of a chimney, is identified by a recognition algorithm which, together with the projection provided by the Ordnance Survey, allows us to geolocate each chimney. An identifier, e.g., 00006, follows the sign. Industry-specific information can then be retrieved after the recognition algorithm has (i) located a chimney and (ii) stored the associated identifier. For instance, the chimney 00006 belongs to *Eastbrook Dye Works (Worsted)*.

We exploit the information on the industrial site associated with each chimney in order to account for sectoral differences in coal use.⁹ We define the following industrial categories: *Brick factories*, *Foundries*, *Chemical factories*, *Mining*, *Breweries*, *Tanneries*, *Food processing*, *Textile production*, *Paper production*, *Shipbuilding* and *Wood processing*. We then run a simple textual analysis based on few keywords for each industry. The algorithm allows to match 90% of the 5,000 chimneys; we classify the remaining 10% under a generic category (*Other manufactures*). We match these categories with aggregate measures of industry-specific coal use per worker (Hanlon, 2019) and we construct the estimated pollution emission E_i from a chimney in industry i as follows:

$$E_i = \frac{C_i \times L_i}{Ch_i}$$

where C_i is the industry-specific measure of coal use per worker, L_i is total employment in industry i , and Ch_i is the total number of chimneys of type i . Accordingly, L_i/Ch_i gives us the average number of workers per chimney in industry i .¹⁰

Pollution dispersion The previous exercise produces a map of industrial pollution sources. We use the *ADMS 5* dispersion model to construct a map of exposure to toxic air pollutants.¹¹ This model is an augmented version of the Gaussian air pollutant dispersion equation known as the Gaussian-Plume model. *ADMS 5* models atmospheric dispersion under a large spectrum of meteorological conditions, provides pollution estimates in coastal areas and incorporates the impact of temperature and

⁹The industrial category—measured with error—is the only information that we use in order to proxy for the average pollution emission associated to a chimney. In theory, we could also exploit the size of the chimney, the size of the associated workshop or a combination of both. The map quality does not allow for a systematic recognition of either the chimney symbol or the contour of the workshop. The measurement error induced by our (imperfect) modeling of pollution emissions is likely to generate an attenuation bias. We show, however, in a robustness check that the industry weight already contains significant information.

¹⁰We report the estimated pollution emission per industry, E_i , in Appendix Table A1. As chimneys in the unidentified category (*Other manufactures*) are apparently associated with small workshops, we calibrate their weight E_i on the least polluting category (*Wood processing*).

¹¹See <http://www.cerc.co.uk/environmental-software/ADMS-model.html>. Atmospheric dispersion models are additive such that concentration of air pollutants is calculated as the sum of concentrations computed separately from each chimney.

humidity. Since industrial chimneys during the Industrial Revolution were shorter than modern chimneys, pollution dispersion was heavily influenced by surrounding topography. Our pollution model thus accounts for complex terrain and changes in surface roughness.

The *ADMS 5* model requires a number of inputs. First, it uses meteorological information for broad regions. We use contemporary 10-year statistical meteorological data as provided by the Met Office, thereby neglecting small changes in prevailing winds related to climate fluctuations between the 19th century and today. Figure 4 illustrates the meteorological information for two of the four regional models: Northern England and Southern England. The “wind roses” show wind provenance and intensity. As expected, winds blow mostly from the west/south-west. Moreover, we see that wind is less predictable in northern England generating on average more dispersed air pollution measures.¹² Second, the model requires complex terrain data and convective meteorological conditions on land. We use the current terrain height and ruggedness, which affect wind speed and turbulence for cities with high gradients.¹³ Finally, *ADMS 5* requires information on the emission source. Atmospheric dispersion modeling is usually parameterized on current chimneys which are tall, wide and have high exit velocity. By contrast, chimneys in the Industrial Revolution were between 10 and 50 meters tall with the majority being shorter than 25 meters. Moreover, the exit velocity and temperature were also lower than today. To incorporate these characteristics, we set chimney height to 25 meters in the baseline and assume an exit velocity of 4 m/s and an exit temperature of 120 degrees Celsius. To model pollution from residential sources, we assume domestic chimneys to be uniformly distributed within city borders, at a very low altitude, and the *ADMS 5* model is used under the same meteorological and topographic inputs.¹⁴

As an example, Figure 5 displays the industrial sources of pollution for Manchester and Oldham (left panel) and the resulting concentration of sulfur dioxide (right panel). Measures of pollution within cities are non-existent in England and Wales before the mid-twentieth century. To validate our pollution measure, we use a sample of deposits collected in few neighborhoods of Manchester by the First Annual Report of the Sanitary Committee on the Work of the Air Pollution Advisory Board,

¹²Under stable conditions and high chimneys, the wind carries pollution far from the origin source while pollution is most intense at the origin under unstable conditions. Note that our benchmark measure uses an average of these different conditions over the past 10 years.

¹³In Appendix Figure A1, we show the differences in pollutant dispersion implied by topography in a city with high gradients (Oldham). Topography and land cover play little role in flat terrains.

¹⁴In our sample of industrialized cities, the relative contribution of domestic emissions (versus industrial emissions) in explaining the distribution of pollutant concentration within city is low: while industrial coal consumption was very high, but also very concentrated in few neighborhoods, residential coal consumption was equally spread across neighborhoods.

1915. We provide a comparison of our constructed measure with this external source in Figure 6. We observe a large variation across neighborhoods for both measures, illustrating that distance to chimneys, topography and wind directions generate large within-city dispersion in pollution with some neighborhoods reaching alarming concentrations in air pollutants.¹⁵ Reassuringly, the estimated pollution very strongly correlates with the deposit measure.

Another validation exercise of our pollution data brings us back to the example of the peppered moth and the appearance of the darker form (*carbonaria*, see Section 1.1). We exploit a collection of surveys reporting the melanic forms of many species of moths (Cook, 2018). We restrict the sample to 54 surveys conducted before the swift decrease in urban pollution (i.e., between 1965 and 1974). Appendix Figure A2 illustrates the relationship between our measure of historical pollution at each survey site and the share of darker moths. We find higher shares of the darker form (*carbonaria*) in highly-polluted areas. There is a difference of about 60 percentage points between the least and most polluted survey sites. This exercise provides additional support for the validity of our pollution measure.

Measure of neighborhood composition Data are drawn from censuses in 1881, 1971, 1981, 1991, 2001 and 2011, and a quasi-census in 1817.

Measures of neighborhood sorting in the nineteenth century are extracted from individual records of the 1881 census which hold information on the structure of households and its members’ gender, age, occupation and place of birth. There are two indicators of household location: a geolocated parish variable and an unreferenced address. While the parish variable is consistently referenced, the address is inconsistently reported (surveyors use abbreviations and misspelling is frequent) and poorly digitized (e.g., due to handwriting). To process this patchy information, we develop a method to allocate households interviewed in the 1881 census to small administrative units (2001 Lower Super Output Area, LSOA), based on the organization of data collection.

On census day, individual surveyors were given blocks to survey and they filled in enumerator books while visiting their allocated neighborhoods: There is a mechanical spatial clustering among adjacent individual records. The exact position

¹⁵To better understand the extent to which cities were polluted at the end of the nineteenth century, we provide the cumulative distribution for our measure of pollution in our sample of neighborhoods (Lower Super Output Areas). Appendix Figure A3 shows that about 10% of our neighborhoods display air pollution above the two National Ambient Air Quality Standards (SO_2 concentration above 12 and $15\mu g/m^3$). About 2% of our sample LSOAs—mostly in Manchester, Oldham and Liverpool—have indices of pollution above the peaks recorded in contemporary Beijing ($40\mu g/m^3$).

of each entry in the 1881 census is an exceptional source of information that has, to the best of our knowledge, not been exploited so far.¹⁶ This position is captured by a book number, a folio number, and a page number. If we are able to locate a fraction of households, we can infer the location of unmatched entries given their position in the census books and the geo-references of well-matched neighbors.

To implement our clustering analysis, we need to geolocate a non-negligible fraction of households. For this purpose, we clean historical addresses and create a pool of geolocated addresses, heritage sites and listed buildings. We then run a fuzzy matching procedure between the pool of census addresses and the pool of geolocated addresses within the same registration parish. We achieve a perfect match for 20% of the total sample, and we match 30% of the total sample with sufficient precision (90% of the original string is found in the matched address).¹⁷

A detailed description of the clustering algorithm is provided in Appendix D and we only discuss its main steps here. First, we define a *cluster id* based on the book, page, and folio numbers for each record. This *id* relates a census entry to its census neighbors. Second, we focus on the sample of well-matched households within each *cluster id*, identify the major cluster of located addresses, isolate its centroid, and associate the overlapping LSOA to this centroid. Third, we attribute this geographic unit to entries with the same *cluster id*, including entries that were not matched during the fuzzy matching procedure. We repeat this algorithm with different cluster definitions, and select the most frequent LSOA identifier. An internal validation of this procedure is performed by randomly excluding well-matched households from the second step; these households are then used as a validation sample where we can compare the attributed geo-location with the actual one.

We use data from the project “The Occupational Structure of England and Wales, c.1817–1881” (Shaw-Taylor and Wrigley, 2014) which, among others, cleans baptism records over 1813–1820 to reconstruct a quasi-census of male occupations around 1817. Individuals can be linked to 834 parishes as defined in the 1881 micro-census. For recent census waves (1971–2011), we use area-weights to map census enumeration districts into 2001 Lower Super Output Area (LSOA). The census data provide consistent measures of occupation, housing, education level and country of origin for all these years.

One drawback is that we do not directly observe income, arguably the best proxy

¹⁶Logan and Parman (2017) exploit the structure of the 1880 U.S. census enumeration to create segregation measures based on the race of “census neighbors”.

¹⁷There are three potential sources of noise when matching historical address with current addresses: (i) reporting error from past surveyors, (ii) digitizing errors and (iii) changes in street names, e.g., red-light districts. The first two sources of error are the most common.

for the social composition of neighborhoods within cities. Instead, we observe 3-digit occupational information in recent censuses and rely on a similar classification for 1817 and 1881 (PST system of classifying occupations; see [Wrigley, 2010](#)). There are various ways to proxy for income based on occupational structure, for example by predicting income using average occupational wages. Such inference would require assumptions regarding the relative wage per occupation across cities. For the sake of transparency, we rely on a proxy based on the raw data, i.e., the share of low-skilled workers among the working population.

For 1817 and 1881, we collapse the 500 occupational sub-categories into 10 categories. We define low-skilled workers as Unemployed, Disabled, Unskilled and Semi-Skilled workers. Managers, Gentlemen, Rentiers, Clerks, and Manual Skilled workers are classified as high-skilled workers. Finally, we assign Farmers to a separate category and drop Soldiers from our analysis. In order to refine our measure, we restrict our sample to individuals with the lowest possible measurement error, i.e., males between 25 and 55.¹⁸ This decomposition brings about 60% of low-skills, 30% of high-skills and 10% of farmers in 1881 (78% of low-skills, 12% of high-skills and 10% of farmers in 1817) in the 70 metropolitan areas.

For 1971–2011, occupations are already classified into 1-digit occupational categories: Managers; Professionals; Associate Professionals; Administration; Manual Skilled; Care; Sales; Processing; and, Elementary. We group the first three categories as high-skills and the remaining six as low-skills to harmonize shares of low-skills between 1881 and 1971–2011. Clerks and Manual Skilled workers are thus classified as low-skills, which brings about 62% of low-skills and 38% of high-skills in 2011. We drop the category “Farmers” as it is close to non-existent among our urban LSOAs.

1.3 Descriptive statistics

We start by providing suggestive evidence about the correlation between exposure to air pollutants and neighborhood composition, and the underlying role of prevailing wind patterns. Figure 7 displays the spatial gradients in pollution and in the share of low-skilled workers at the end of the nineteenth century in the average city of our sample.¹⁹ The pollution cloud unambiguously leans towards the east, due to

¹⁸Our results are robust to (i) adding female workers as we will show later, and (ii) widening the age interval (e.g., 15–65).

¹⁹The left panel of Figure 7 is constructed as follows. First, we define a grid of equally-spaced points—every 100 meters—within 1.5 kilometers of each city centroid, and we associate excess pollution at each point relative to the average city pollution. Second, we overlay the city grids and compute, for each point, the unweighted average of excess pollution across cities. Third, we interpolate across grid points using a Gaussian Kernel interpolation method. The 20 level lines are quantiles of pollution. The right panel is constructed in a same fashion with the share of low-skilled

prevailing winds or possibly the unequal distribution of pollution sources across space. The spatial gradient in the share of low-skilled workers also exhibits a similar asymmetric pattern towards eastern neighborhoods, albeit slightly noisier.

We refine our analysis of the relationship between the share of low-skilled workers in 1881 and pollution sources in Figure 8. Units of observation are the 675,000 block \times chimney pairs where a block is a census cluster of households with the same geolocation in 1881 which is located within 2 kilometers of the chimney (there are 100,000 such blocks). The left panel of Figure 8 displays the average share of low-skilled workers in 1881, as a function of distance to the chimney. There is a sharp gradient, with a 10 percentage point difference between 100 and 1,500 meters from a pollution source. This gradient likely captures high commuting costs. Strikingly, even conditional on distance to the pollution source and amenities in 1881, there remains large variation in the share of low-skilled workers at the block level. Part of this variation relates to the location of the block relative to the chimney. As apparent in the right panel of Figure 8, there is a 1.5 percentage point excess share of low-skilled workers for blocks situated north-east of the chimney. This gradient in the direction of prevailing winds is what will be captured in the reduced-form analysis of the next section.

Table 1 provides summary statistics at the level of our baseline units of observation. Within a buffer of 20 kilometers around the centroids of our 70 cities, the clustering process associates about 5 million active male workers in 1881 to 5,538 LSOAs. As these LSOAs are 2001 census units, we can associate contemporary measures to all of these 5,538 geographic units which will constitute our baseline sample. We provide summary statistics for the full sample, and for LSOAs with above- and below-median pollution at the city level. We report statistics for the main outcome variables and baseline controls accounting for topography, amenities, and direction (latitude and longitude). Some of these characteristics capture important differences between more and less polluted LSOAs within cities. Less polluted neighborhoods have higher elevation, are more rural, and are more distant from waterways and pollution sources. Interestingly, they are on average 2,000 meters further to the West than highly polluted neighborhoods. In the last three columns of Table 1, we provide a decomposition of the variance within and between cities. A very large share of the variance in pollution is within cities. Our empirical strategy, described in the following section, hinges on such within-city variation and is mostly orthogonal to variation across cities.

workers in 1881. We provide a similar illustration centered on town halls in Appendix Figure A4.

2 Reduced-form evidence

This section presents reduced-form evidence on (i) historical pollution and neighborhood sorting, and (ii) the subsequent persistence of neighborhood segregation.

2.1 Empirical strategy

To estimate the impact of pollution on neighborhood sorting within cities, we run a baseline difference specification at the LSOA level and an IV specification where we employ exogenous factors that influenced the location of pollution sources.

Baseline specification Letting i denote a LSOA, p a parish, c a city, and t a particular census wave, we estimate the following equation:

$$Y_{it} = \alpha + \beta P_i + \gamma \mathbf{X}_i + \nu \mathbf{Y}_p + \delta_c + \varepsilon_{ict} \quad (\text{S1})$$

where Y_{it} is a measure of occupational structure. The measure of historical pollution, P_i , results from a combination of the location of pollution sources and a dispersion process. Physical features like hills or rivers that enter the simulated pollution measure may also be local (dis)amenities that affect individual neighborhood choices. To eliminate this potential source of bias, we include separate topography indicators (e.g., maximum, minimum and average elevation) along with a rich set of geographic controls (e.g., area, share of LSOA within the city borders, latitude and longitude) and controls for (dis)amenities (distance to waterways, heavy-industry chimneys, light-industry chimneys, the town hall and parks) in the set of controls \mathbf{X}_i . \mathbf{Y}_p is a set of measures of occupational structure in 1817 at the parish level (shares of low-skills, high-skills and farmers) and the logarithm of the property tax in 1815 at the parish level, which capture possible fixed neighborhood amenities. δ_c are city fixed effects and standard errors are clustered at the parish level.²⁰

We further exploit the interaction between the distribution of pollution sources and air pollution dispersion by considering counterfactual diffusion processes, e.g., as generated by the same pollution sources but with rotated wind patterns. One can think of this procedure as a decomposition of the interaction between location and diffusion; it isolates variation induced by the asymmetry between neighborhoods at the same distance from factories, some of them being located downwind and others

²⁰There may exist sampling variation in our main independent variable, e.g., as induced by the measured average weather conditions, because it is a simulated measure of pollution. This sampling variation occurs within the atmospheric dispersion model; we thus cannot correct standard errors by bootstrapping.

upwind (as in Figure 8).

A concern with specification (S1) is that the treatment may not be exogenous because fixed unobserved amenities explain both the upwind presence of industries and the local occupational structure. In robustness checks, we show a balance test before the rise of industrial coal pollution, and provide identification at a more granular level by including fixed effects at the level of parishes or electoral wards.²¹

Finally, there is a remaining threat to identification from reverse causality or time-varying omitted variation. For instance, factories may be strategically placed upwind of poor neighborhoods to minimize political or economic costs associated with environmental disamenities in richer neighborhoods. We address this concern with an instrumental variable.

IV specification To account for the bias arising from non-random industry location, we exploit exogenous variation in location factors which translates into exogenous variation in pollution imprints. Specifically, we exploit the fact that large boilers required a constant stream of water for cooling. As a result, the natural geographic placement of all mills was along rivers or canals (Maw et al., 2012). We locate hypothetical chimneys in intervals of 150 meters along waterways in 1827, before the rise of coal as the main energy source. To derive our instrument, we assume uniform air pollutant emissions from these exogenous pollution sources, combined with the actual atmospheric dispersion due to wind flows and topography.²² This natural geographic placement of chimneys is not susceptible to being selectively placed upwind of poor neighborhoods. However, the variation correlates with proximity to waterways which may itself affect the attractiveness of a neighborhood. We thus control separately for distance to waterways and exclude neighborhoods bordering these waterways.

We then use the following first-stage specification to instrument the historical pollution, P_i , in Equation (S1):

$$P_i = b_0 + b_1 PP_i + \mathbf{cX}_i + d_c + \mathbf{fY}_p + e_{ict} \quad (\text{S2})$$

²¹One issue with our reduced-form approach is that the potential outcome for one neighborhood is affected by the treatment intensity in other neighborhoods. Finer fixed effects may aggravate this issue if individuals choose residences within a small radius around their working place. One way to deal with this issue is to develop a proper model of neighborhood choice which accounts for equilibrium adjustments at the city-level (see Section 3).

²²Appendix Figure A5 describes our approach. In panel (a), we see the cities of Manchester and Oldham with the associated 1827 natural waterways. Panel (b) displays the natural geographic placement of chimneys along canals and panel (c) the resulting spatial distribution of air pollutants. Finally, panel (d) shows the distribution of air pollutants using actual pollution sources.

where PP_i is the simulated pollution using hypothetical chimneys. As described above, \mathbf{X}_i includes a comprehensive set of controls for physical attributes, \mathbf{Y}_p is the occupational structure in 1817 at the parish level and $\{d_c\}$ are city fixed effects.

2.2 Historical pollution and neighborhood sorting

In this section, we document a negative correlation between air pollution and neighborhood income, as proxied by the share of low-skilled workers in 1881.

In Table 2, we report the estimates for our baseline specification (S1). As can be seen in the first column, air pollution and the share of low-skilled workers in 1881 are positively correlated. Controlling for a large set of covariates does not affect the estimates. In the second column, we add city fixed-effects to control for variation in atmospheric pollution and neighborhood composition between cities (Hanlon, 2019). In the third column, we add (log) property tax in 1815, and the parish-level shares of low-skilled workers, high-skilled workers and farmers in 1817 to clean for potentially unobserved fixed characteristics. From the fourth to the last column, we add separate elements entering in the pollution dispersion process. In the fourth column, we condition on topography (elevation and distance to waterways in 1827). In the fifth column, we control for distance to pollution sources (heavy- and light-industry), distance to the city hall, distance to parks, area and the share of the LSOA within the 1880 city borders.²³ In the sixth column, we add eastings and northings of the LSOA centroids to control for wind patterns and potential western or southern preferences in locations. As apparent from Table 2, our estimates slightly decrease but remain large and precisely estimated.²⁴

The correlation between air pollution and the occupational structure is both statistically and economically significant. In the baseline specification (column 6), the coefficient is 0.033 and the 95%-confidence interval is [0.020, 0.047]. One additional standard deviation in air pollution increases the prevalence of low-skilled workers by 3.3 percentage points, which is slightly less than 15% of a standard deviation in their prevalence across LSOAs. A differential in pollution equivalent to the one between the first and last deciles in Manchester would be associated with a differential of 16 percentage points in the share of low-skilled workers.²⁵ Residential sorting may

²³As stated in Section 1, the metropolitan areas (and thus the sample of LSOAs) are defined by a buffer of 20 kilometers around the centroids of our cities. In robustness checks, we verify that the results are left unchanged if we limit the sample to urban LSOAs intersecting with the 1880 city borders (which may be endogenous and affected by pollution through agricultural yields).

²⁴The first column of Appendix Table A3 reports the coefficients on all covariates.

²⁵We consider other outcomes in Appendix Table A2, with the share of all low-skilled workers including females and the share of migrants, distinguishing between migrants from England and Wales and from the Commonwealth. We find that the standardized effects of pollution on the

be tempered by the necessity for residents to live close to their working place, as induced by relatively costly modes of transportation.

Figure 9 illustrates the estimated relationship between the share of low-skilled workers before and after the rise in coal use (respectively in 1817 and in 1881) and the atmospheric pollution during the Industrial Revolution. On the y -axis, we plot the residuals from a regression of the standardized shares of low-skilled workers on a similar set of controls as in column 6 of Table 2. On the x -axis, we plot the regression-adjusted residual of standardized air pollution. The relationship between the share of low-skilled workers and standardized air pollution is strongly positive but flattens at both extremes, i.e., for very high and very low within-city pollution levels. By contrast, there is no correlation between the share of low-skilled workers in 1817 and the measure of atmospheric pollution.

The most convincing evidence in support of a causal relationship between pollution and neighborhood sorting comes from hypothetical pollution imprints. We generate pollution exposure with existing chimneys and a *rotated* wind profile in steps of 30 degrees around each source, and we estimate its conditional correlation with the share of low-skilled workers. In this exercise, we control for a measure of *static pollution* capturing proximity to the pollution source. We construct this measure from existing chimneys under a static wind profile that is symmetric in all directions. Figure 10 illustrates the results for the years 1817 and 1881 where we center the figure around the actual pollution pattern at 0 degrees. In 1817 (Panel a), *before* the rise of coal pollution, we observe virtually no correlation between the pollution measure at any degree and the share of low-skilled workers. In 1881 (Panel b), *after* pollution became a relevant disamenity, we observe a pronounced, bell-shaped pattern with the peak around the actual pollution pattern between 0 and 30 degrees.²⁶ As we rotate prevailing winds, the estimated relationship loses significance and turns negative. To reduce measurement error, we clean our estimates for parish fixed effects in Panel (c) and for residential pollution in Panel (d). The estimates remain large within a narrow corridor along prevailing winds, but they now decrease sharply, becoming negative for rotations of more than 90 degrees.

Finally, we present in Table 3 the results of the IV strategy (S2), that uses 1827

share of all low-skilled workers and migrants are comparable to the baseline findings. Interestingly, the higher prevalence of migrants in polluted neighborhoods is essentially due to migrants from England and Wales (thus unrelated to the Irish Potato Famine).

²⁶The fact that the peak in correlation is between 0 and 30 degrees may be due to measurement error. First, wind patterns may have changed in one century, in particular the frequency of cyclonic or anti-cyclonic conditions (Lamb, 1972), each associated with different wind direction profiles. Second, we consider yearly averages for our meteorological conditions, possibly ignoring differential pollution exposure and wind patterns across seasons or hours of a day.

waterways as a source of exogenous variation for chimney location. We report four sets of estimates, excluding neighborhoods within 250m or 500m of waterways, and with or without extended controls in addition to city fixed-effects. The first stage is strong in all cases, and the (stable) 2SLS estimates tend to be larger than the OLS estimates. One additional standard deviation in air pollution increases the prevalence of low-skilled workers by about 9 percentage points. One explanation for the downward bias in the OLS specification could be measurement error or possibly the difference between the local average treatment—relying on small variations in possibly highly-polluted neighborhoods—and the average treatment effects.

2.3 Historical pollution and contemporary neighborhood segregation

This section expands our previous analysis of neighborhood sorting to recent census waves (1971–2011) to assess potential reversion to the mean after the 1968 Clean Air Act stopped coal use within cities. Table 4 reports the slopes between the shares of low-skilled workers and historical pollution, as estimated by Equation (S1). One standard deviation in historical air pollution increases the prevalence of low-skilled workers by 2.5 to 4 percentage points without a clear pattern between 1971 and 2011, and the standardized effects range between 0.19 and 0.23.²⁷

Our analysis of spatial inequalities in cities of the nineteenth century was not informed by house or land prices. In the recent period, however, we do observe possible differences in housing demand across neighborhoods through transaction prices. In Table 6, we use transactions in England and Wales as recorded by Land Registry between 2000 and 2011 (columns 1 and 2) and by Nationwide Building Society between 2009 and 2013 (columns 3 and 4), and we run hedonic regressions with and without controlling for average house characteristics. We find that one additional standard deviation in past pollution is associated with a price drop of about 10% to 11%. Controlling for property characteristics reduces these estimates to 6-8%, showing that properties in formerly polluted neighborhoods are smaller and more likely to be “non-detached”. Figure 11 illustrates these very large effects.

Past environmental disamenities appear to have a marked effect on spatial inequalities today. A differential in pollution equivalent to the one between the first and last deciles in Manchester is still associated with a differential of 16 percentage points in the share of low-skilled workers or with differences in property prices of about 40%, thereby explaining the social gradient between neighborhoods in the

²⁷Appendix Table 5 displays the IV estimates for the occupational structure in recent years, and shows, as in Table 4, that there are no signs of reversion to the mean.

west and east often evoked in popular culture.²⁸

To visualize possible non-linearities in the persistence of neighborhood sorting, Figure 12 displays the relationship between shares of low-skilled workers in 1817 (long dash), 1881 (short dash), 1971, 1991 and 2011 (plain lines) and the historical pollution disamenity that stopped after the 1968 Clean Air Act. As apparent, we observe some reversion to the mean for low and intermediate values of within-city pollution. By contrast, segregation patterns appear to persist at around one standard deviation above average within-city pollution.

To take a closer look at the underlying dynamics, we refine the analysis between 1971 and 2011 and organize the data in a panel structure with decadal observations for each LSOA. To shed light on non-linearities, we define ten pollution categories corresponding to the ten deciles in intra-city pollution, i.e., neighborhood pollution adjusted by the average city pollution. We then run a panel regression with the share of low skilled workers as the dependent variable, LSOA fixed-effects, city \times year fixed effects and trends for each pollution decile. The initial relationship between the share of low skilled workers in 1971 and pollution deciles is displayed in the left panel of Figure 13, while the estimates for pollution-decile trends are shown in the right panel. A process of mean-reversion would be captured in the form of a decreasing pattern in trends for each pollution decile.²⁹

Our findings are not consistent with a uniform reversion to the mean. The right panel of Figure 13 shows evidence of mean-reversion over the period 1971–2011 but only for areas with below-median levels of past pollution exposure (i.e., pollution categories 1–5). We see a slight increase in the share of low-income workers in the least polluted neighborhoods and a corresponding decrease in moderately polluted neighborhoods. For neighborhoods with above-median levels of past pollution, however, we see the opposite. These neighborhoods become even more deprived relative to the median polluted neighborhoods. This pattern would be consistent with tipping dynamics leading to a high persistence of deprivation in neighborhoods with extreme pollution exposure in the past.

²⁸The persistence of the relationship between historical pollution and neighborhood composition cannot be mechanically attributed to the collapse of industries in the former *cottonopolis*, given that our estimates are identified within cities.

²⁹Letting $\{x_{i,t}\}$ denote the excess share of low-skilled workers in neighborhood i and period t and $\theta < 1$ the AR(1) parameter, we should observe that:

$$E[x_{i,t+1} - x_{i,t}] = (\theta - 1)E[x_{i,t}]$$

2.4 Sensitivity and robustness checks

We conduct a large number of robustness checks around the baseline specification(s). We summarize the findings in this subsection and leave a detailed discussion of the results along with additional Figures and Tables to Appendix E.

First, we conduct balance tests in the period before coal pollution to further reduce concerns about biasing effects from unobserved pre-existing neighborhood characteristics. There is no correlation at the parish level between either the 1817 share of low-skilled workers or the 1815 property tax returns (as a proxy for wealth) and the later atmospheric pollution.

Second, we consider variations in our pollution modeling. We only vary the chimney height as the exit velocity and the exit temperature would affect the same crucial model input, i.e., the height of the smoke column in the atmosphere. The atmospheric pollution based on chimneys assumed to be shorter (15m) or taller (40m) than our baseline (25m) generates similar estimates as in Tables 2 and 4.

We also use a simple, albeit less informative, measure of pollution exposure. We overlay each city with a grid of equally-spaced points and count the number of chimneys that are located in (i) the North-East, (ii) the North-West, (iii) the South-East, or (iv) the South-West quadrants—within a given distance from each grid point. Lastly, we collapse the measure for all four quadrants into 2001 LSOAs. Only chimneys located North-West and South-West are correlated with the share of low-skilled workers, which coincides with prevailing downwind directions. Once we split the sample into Northern and Southern cities, we find that chimneys located in South-West quadrants are relatively more predictive of deprivation in Northern cities, reflecting the more southerly wind direction in North England (see Figure 4).

Third, we test the sensitivity of our findings to the addition of the following pollution imprints as controls: a ‘static’ placebo pollution measure which assumes that pollution spreads evenly in all directions and captures the proximity to industrial centers; a placebo pollution pattern which varies the emission intensity by coding the chimneys of high (low) polluting industries as low (high) polluting; residential pollution; and, a contemporary measure of atmospheric pollution. We find that the two placebo patterns do not affect the significance of the main pollution variable and have no predicting power, suggesting that there is no additional information in distance or emission intensity that might affect our estimation. Residential pollution does predict deprivation, but its standardized effect is one third of the standardized effect of industrial pollution. Contemporary pollution has a relatively small impact on neighborhood composition in 2011. Historical pollution is far more predictive of current spatial inequalities than current exposure to air pollutants.

Fourth, we explore sensitivity to controls, fixed effects, clustering and sample selection. We control for a number of within-city geography variables interacted with city-fixed effects to capture city-specific geographic patterns and commuting infrastructure: latitude and longitude; distance to the town hall; distance to heavy industries; and, distance to light industries. The effect of a standard deviation in pollution remains stable at between 3.6 and 4.5 percentage points in the share of low-skilled workers. We also report the results of our baseline specification with 540 parish-fixed effects, with fixed effects expanded to 1,440 electoral wards and to the 1,850 Middle Layer Super Output Areas. The estimates remain unchanged, even when identification comes from a within-MSOA comparison. We report standard errors clustered at three different levels: electoral ward; MSOA; and, city. Standard errors increase by about 40% between the least and most conservative choice, and our baseline analysis clustered at the parish-level is at the center of this interval. Finally, we estimate the baseline specification on alternative samples where we exclude: Greater London; the North-West; the North-East. The estimates fluctuate around the baseline, but they remain large in all cases. We also analyze the sensitivity of our results to the exclusion of suburbs and rural LSOAs. The estimates remain precisely estimated and, in some cases, slightly larger than in the baseline.

Fifth, we consider an alternative instrument, based on the historical settlements of industries. We isolate variation induced by the location of industrial districts *before* coal became a major energy source which would affect disproportionately downwind neighborhoods. To predict early industrial districts, we locate 543 early steam engines installed between 1700–1800, using data from [Kanefsky and Robey \(1980\)](#) and [Nuvolari et al. \(2011\)](#). We model uniform air pollutant emissions from early steam engine locations and use the resulting atmospheric dispersion as an instrument for actual pollution, conditioning for distance to the nearest industrial chimneys. The estimates from using this instrument are remarkably similar to the ones using waterways as an exogenous factor for industry location.

Finally, we provide some insight on the recent dynamics of neighborhood composition (1971–2011) in Appendix [F](#) where we document differences across neighborhoods in school supply, crime, housing quality and public amenities. We also show how the liberalization of social housing and immigration inflows may have contributed to residential segregation between 1971 and 2011 in Appendix [G](#).

3 A dynamic model of residential sorting

In order to quantify the non-linear dynamics in the persistence of neighborhood sorting between 1971 and 2011, we develop a dynamic model of residential sorting

within a city, in which infinitely-lived households face relocation frictions. Those relocation frictions mean that optimal residential choice depends on past and future neighborhood amenities. The purpose of the model is to derive a dynamic demand equation for neighborhoods, which we can then proceed to estimate in Section 4.

Environment Consider a unit mass of infinitely-lived households, each of measure zero. Time is discrete and each household receives utility every period from the consumption of the numeraire and a neighborhood amenity. Households are hand-to-mouth consumers, and they rent land from absentee landlords in a closed city. Land markets are competitive and land supply is inelastic and constant over time.³⁰

We assume that there is a discrete number of neighborhoods J . Let $r_{j,t}$ denote the rental cost in neighborhood j and period t . The amenity in each neighborhood may be time-varying and we denote it with $a_{j,t}$. Finally, we assume that there is a household- and period-specific idiosyncratic preference shock, $\varepsilon_{i,j,t}$, for household i in neighborhood j and period t .

The flow of utility for household i residing in neighborhood j and period t is $u_{j,t} + \varepsilon_{i,j,t}$, where $u_{j,t}$ depends on consumption and the amenity as follows:

$$u_{j,t} = g(a_{j,t}, y_t - r_{j,t}),$$

and y_t is the (exogenous) income in period t .

At the beginning of each period t , the idiosyncratic preference shock is revealed and there is another household-specific idiosyncratic draw: with probability $1 - \theta$, the household can freely relocate to any other neighborhood within the period. This relocation shock is a convenient way to capture the presence of moving rigidities. This formalization has two important properties. First, there is an exogenous and representative share θ of non-movers in each period. Second, and in contrast with implications of the common assumption of fixed moving costs (Bayer et al., 2016), the location choice of a possible mover is not tied to their previous location.³¹ One interpretation is that the psychological cost of considering relocation is very high, and only an “external” event can force the household to pay such a cost. For instance, a liquidity shock may force the absentee landlord to sell the property—the equivalent of exogenous firing in models of labor search. Alternatively, the household may be affected by life-cycle shocks (e.g., the birth of a child).

³⁰The assumption that land supply is inelastic and constant over time is not crucial to derive the dynamic demand for neighborhoods, but is required to identify the model with our data.

³¹The *possibility* to relocate is orthogonal to the current location and to the relative valuations of the different neighborhoods in the city. Conditional on being able to move, however, the household decision to stay or move will account for relative preferences for neighborhoods.

Letting β denote the discount factor, the value of residing in neighborhood j and period t for household i is:

$$U_{i,j,t} = u_{j,t} + \varepsilon_{i,j,t} + \beta\theta E_t U_{i,j,t+1} + \beta(1-\theta) E_t V_{t+1},$$

where V_{t+1} is the value function for a household with the opportunity to relocate at $t+1$. A household with the opportunity to relocate considers the path of future (and possibly stochastic) amenities $\{a_{j,t}\}_{j,t}$ as given, and maximizes:

$$V_t = \max_j \{u_{j,t} + \varepsilon_{i,j,t} + \beta\theta E_t U_{i,j,t+1}\} + \beta(1-\theta) E_t V_{t+1}$$

We assume, as common in models of residential sorting (Bayer et al., 2016), that idiosyncratic preferences, $\varepsilon_{i,j,t}$, are distributed across households along a Type 1 Extreme-Value distribution such that the share of households opting for neighborhood j if they have such option in period t , $n_{j,t}^*$, follows a logit model:

$$n_{j,t}^* = \frac{e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}}}{\sum_j e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}}}. \quad (1)$$

Note that households only need to consider the states of nature along which no further re-optimization is possible; these occurrences are the only ones in which the current relocation decision matters.

In equilibrium, rental prices, $r_{j,t}$, adjust such as to equate land supply and land demand in each neighborhood.

Equilibrium with two types of households Consider that the city is populated by two types of households $o \in \{H, L\}$ in proportion $(\eta, 1-\eta)$, differing only along their income, where $y^L < y^H$. The demand for neighborhood j is given by Equation (1),

$$\begin{cases} \ln(n_{j,t}^{H*}) = \sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^H - \ln\left(\sum_j e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^H}\right) \\ \ln(n_{j,t}^{L*}) = \sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^L - \ln\left(\sum_j e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^L}\right) \end{cases}$$

where $(n_{j,t}^{H*}, n_{j,t}^{L*})$ are the fractions of each type of household opting for neighborhood j at time t . Letting $(n_{j,t}^H, n_{j,t}^L)$ denote the fractions of each type of household *residing* in neighborhood j at time t , we have that:

$$\begin{cases} n_{j,t}^H = \theta n_{j,t-1}^H + (1-\theta)n_{j,t}^{H*} \\ n_{j,t}^L = \theta n_{j,t-1}^L + (1-\theta)n_{j,t}^{L*} \end{cases}$$

and the land market equilibrium implies that, in each period t and neighborhood j ,

$$\eta n_{j,t}^H + (1 - \eta) n_{j,t}^L = n_j,$$

where n_j is the fixed land supply in neighborhood j . We can use these equilibrium conditions in periods t and $t - 1$ in order to express the whole problem as a function of the *share* $s_{j,t}$ of type- L households among households living in neighborhood j and period t .³² By subtracting the respective demand schedules, we obtain:³³

$$F\left(\frac{s_{j,t} - \theta s_{j,t-1}}{1 - \theta}\right) = \sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t(u_{j,t+\tau}^H - u_{j,t+\tau}^L) + \mu_t.$$

We can thus write the dynamic demand equation for neighborhood j as:

$$F\left(\frac{s_{j,t} - \theta s_{j,t-1}}{1 - \theta}\right) - \beta\theta E_t F\left(\frac{s_{j,t+1} - \theta s_{j,t}}{1 - \theta}\right) = u_{j,t}^H - u_{j,t}^L + \nu_t, \quad (2)$$

where F is a decreasing function, and $\nu_t = \mu_t - E_t \mu_{t+1}$ is a time-fluctuating variable which captures dynamics in the relative welfare of type- L residents and does not vary across neighborhoods.

Equation (2) characterizes the relative demand for neighborhood j . The right-hand side is the contemporaneous valuation of living in neighborhood j for an average type- H household relative to an average type- L household. The left-hand side is the relative demand for this neighborhood. A positive shock to the relative valuation of neighborhood j at time t , $u_{j,t}^H - u_{j,t}^L$, induces an instantaneous decrease in the share of low-skilled workers residing in neighborhood j . This decrease is however tempered by relocation rigidities: (i) some residents have not been able to respond to the change in valuation by relocating, (ii) the residents that do have the opportunity to relocate account for the fact they may not be able to relocate in subsequent periods.

³²The reader interested in the derivation of the demand schedule can refer to Appendix B where we also study (i) the dynamic properties of the system of equations, and (ii) an extension with many household types and elastic land supply. Note that we normalize the total number of households to be equal to 1 without loss of generality, i.e.,

$$\eta \sum_j n_{j,t}^H + (1 - \eta) \sum_j n_{j,t}^L = 1.$$

³³We define $F(x) = \ln(1 - x) - \ln(x)$ and

$$\mu_t = \ln\left(\sum_j e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^L}\right) - \ln\left(\sum_j e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^H}\right) + \ln\left(\frac{\eta}{1 - \eta}\right).$$

The quantity μ_t captures the average relative welfare of each type in period t .

With $\theta = 0$, the program of residents collapses to a static problem. They do not need to worry—in the current period—about location in subsequent periods, and they fully adjust to any changes in their valuation of neighborhoods. With $\theta > 0$, the implied demand depends on the past and future expected allocation of households. The next section discusses the identification and estimation of Equation (2) in the data.

4 The persistence of residential sorting

In this section, we discuss the identification of the relative demand for a neighborhood in the data, present the estimates and discuss counterfactual exercises.

Identification We estimate Equation (2) using the observed neighborhood composition in 70 closed cities, indexed by c , over five waves (1971, 1981, 1991, 2001, 2011). The identification requires the following assumptions. We suppose that utility features complementarity between the neighborhood amenity, $a_{j,c,t}$, and the consumption of the numeraire, and that household types only differ along their fixed income y . Specifically, we let the utility of a type- o household living in neighborhood j of city c be,

$$u_{j,c,t}^o = a_{j,c,t} (y^o - r_{j,c,t}).$$

We further assume that the neighborhood amenity, $a_{j,c,t}$, is a function of time-invariant neighborhood characteristics, $a_{j,c}$, and an endogenous amenity proxied by the current composition of the neighborhood, $s_{j,c,t}$.³⁴ We show in Appendix B that demand for neighborhood j can then be written as,

$$F\left(\frac{s_{j,c,t} - \theta s_{j,c,t-1}}{1 - \theta}\right) - \beta \theta E_t F\left(\frac{s_{j,c,t+1} - \theta s_{j,c,t}}{1 - \theta}\right) = h(a_{j,c}, s_{j,c,t}) + \nu_{c,t}.$$

where the function h is a linear transformation of the neighborhood amenity, $a_{j,c,t}$. The estimation of this equation still requires specifying a functional form for the relative valuation of neighborhood j in city c . Our data do not include resident flows between two neighborhoods (as in Bayer et al., 2016, for instance), even at the aggregate level. Therefore, we can only identify a function which is separable in fixed unobserved neighborhood characteristics $b_{j,c}$, observed neighborhood charac-

³⁴The composition of the neighborhood is used as a proxy for the wide range of neighborhood effects which may affect the perceived “neighborhood quality”. Durlauf (2004) and Rosenthal and Ross (2015) are excellent overviews of neighborhood effects affecting residential choices. These effects may include school quality (Durlauf, 1996), quality of the housing stock (Rosenthal, 2008), preferences to live among workers of similar or higher income groups (Guerrieri et al., 2013) or the same ethnic group (Card et al., 2008).

teristics and an unobserved noise $\varepsilon_{j,c,t}$. Intuitively, fixed unobserved neighborhood characteristics can be captured by neighborhood fixed-effects, and the only time-varying unobserved heterogeneity that could be permitted would be at the city level (and then captured by city \times year fixed-effects). We capture possible complementarities between the exogenous and endogenous amenities by interacting a vector of *observable* neighborhood characteristics, $\mathbf{a}_{j,c}$, and $s_{j,c,t}$:

$$F\left(\frac{s_{j,c,t} - \theta s_{j,c,t-1}}{1 - \theta}\right) - \beta \theta E_t F\left(\frac{s_{j,c,t+1} - \theta s_{j,c,t}}{1 - \theta}\right) = \alpha \cdot (\mathbf{a}_{j,c} \times s_{j,c,t}) + b_{j,c} + \varepsilon_{j,c,t} + \nu_{c,t} \quad (\text{S3})$$

In specification (S3), the parameters of interest are θ , β and the vector of parameters characterizing the relative valuation of neighborhoods, α .³⁵ The estimation needs to absorb city/wave-fixed effects, $\{\nu_{c,t}\}$, and neighborhood-fixed effects, $\{b_{j,c}\}$.

We estimate Equation (S3) by Generalized Method of Moments. One concern is that the contemporary share of low-skilled workers, $s_{j,c,t}$, may be correlated with the unobserved noise $\varepsilon_{j,c,t}$. We thus exclude the share of low-skilled workers, $s_{j,c,t}$, from the set of instruments and replace it with the measure of historical pollution interacted with time. Intuitively, this empirical specification uses past pollution as a shifter for the (initial) share of low-skilled workers and the subsequent readjustment in neighborhood composition.³⁶ The identification hypothesis is that historical pollution only affects relative demand for neighborhoods through inherited neighborhood composition; it relies on the observation that industrial pollution from coal-burning disappears from cities after 1971.

Structural estimates The estimates for Equation (S3) are reported in Table 7. The first column reports a specification where neighbourhood preferences are linear in neighborhood composition, the second column adds the interaction with an amenity index predicted by a set of observable geographic characteristics.³⁷

³⁵The system of Equations (S3) admits a unique solution $\{s_{j,c,t}\}$ under some conditions that are made explicit in Appendix B.

³⁶Specification (S3) cleans for neighborhood-fixed effects: an instrument should thus vary over time and predict the *evolution* of neighborhood composition. We rely on the insights from the reduced form analysis that initial pollution does not only shape the initial share of low-skilled workers in 1971 but also its future evolution: we interact the measure of historical pollution with time. The instrument is crucial for identification, as initial neighborhood composition may reflect an unobserved neighborhood amenity which may also directly influence future preferences for the neighborhood. Finally, we control for the interactions of (i) the initial share of social housing in 1971 and (ii) bombing intensity during the German Blitz, with time, in order to clean for dynamics induced by social housing policies (e.g., the Housing Act 1980) and urban renewal. We study the respective roles of social housing and bombing intensity on neighborhood sorting in Appendix G and Appendix H.

³⁷We construct the index by isolating the prediction from a regression of the share of low-skilled workers in 1971 on the average, maximum and minimum elevations for the LSOA, the (inverse)

The share of low-skilled workers lowers the valuation of the neighborhood (see column 1), which generates a multiplier effect. One additional standard deviation in the share of low-skilled workers (about 7 percentage points) reduces the relative valuation of a neighbourhood by about 0.22, which would trigger—off-equilibrium and at the steady-state—a further decrease in the share of high-skilled workers of 3 percentage points. This finding is illustrated in Appendix Figure A7. The persistence in neighborhood dynamics derives from the existence of neighborhood effects, combined with non-negligible moving rigidities. Past neighborhood composition directly influences current neighborhood composition, as many residents will not be given the opportunity to relocate. “Movers” anticipate this long-lasting effect on neighborhood composition and adjust their location choices accordingly. We quantify this indirect, forward-looking effect in the next section.

The probability to be given the opportunity to relocate over a period of ten years, $1 - \theta$, is estimated to be 51%. Assuming that residents always move when given the possibility (they would almost certainly do so in a city with many neighbourhoods), this estimate implies an average housing tenure of about 14 years, consistent with average rates of turnover observed in the housing market. The annualized discount rate, estimated to be around 0.8%, is in the low range of risk-free rates observed over the past 50 years (in spite of a recent drop in such rates, [Del Negro et al., 2019](#)).

How residents value neighborhood composition also depends on exogenous amenities. More specifically, the marginal effect of neighborhood composition on neighborhood valuation decreases with fixed disamenities. This finding could be rationalized by an anchoring effect whereby permanent disamenities affect the desirability of a neighborhood irrespective of its current composition. In essence, the presence of richer residents in a neighborhood increases its desirability, but more so in areas with local exogenous amenities, a mechanism which would be consistent with the (rapid) gentrification process observed in well-connected urban centers ([Guerrieri et al., 2013](#); [Baum-Snow and Hartley, 2016](#); [Couture and Handbury, 2019](#)).

Simulation and counterfactual experiments We simulate neighborhood dynamics between 1991 and 2011 using the structural estimates reported in column 1 of Table 7, the observed fixed amenities—which are used to proxy for the steady-state neighborhood composition—, and the actual shares of low-skilled workers in 1971 and 1981. Panel (a) of Figure 14 shows the model performance in explaining the transitional dynamics in 1991, 2001 and 2011. There is a secular and exoge-

distance to waterways as of 1827, the (inverse) distance to the city hall in 1880, the (inverse) distance to parks in 1880, the share of LSOA within the city borders in 1880, the LSOA area, the latitude and longitude of the LSOA centroid.

nous decrease in the share of low-skilled workers. This decrease is however more pronounced for neighborhoods with low historical pollution exposure, where neighborhood composition does not anchor expectations about future amenities. In effect, the relationship between past pollution and deprivation is at least as strong in 2011 as in 1971, when industrial pollution abruptly waned from urban centers. We report the observed shares of low-skilled workers in 2011 to provide some visual evidence of the model fit: the model reproduces the general neighborhood dynamics but slightly over-estimates the persistence of deprivation in formerly highly-polluted neighborhoods.

We then provide two counterfactual experiments to understand the respective role of preferences and location frictions in the non-linear and persistent dynamics of segregation. In a first counterfactual experiment, we simulate neighborhood dynamics in a model where preferences are set to be orthogonal to current neighborhood composition. In panel (b) of Figure 14, we report the transitional dynamics in the (standardized) share of low-skilled workers in 1991, 2001 and 2011. The secular decrease in the share of low-skilled workers between 1971 and 2011 is much larger in formerly highly-polluted neighborhoods: Neutralizing neighborhood effects is sufficient to turn around the dynamics of segregation, the remaining correlation between past pollution and deprivation being entirely explained by fixed geographic amenities. In a second counterfactual experiment, we keep preferences as in the baseline model and instead modify the extent of rigidities in the relocation process. More precisely, we increase the annual probability to be given the opportunity to relocate from 7% to 20% (corresponding to $\theta = 0.10$) such that almost all residents have the opportunity to relocate between two Census waves. We report the transitional dynamics in panel (c) of Figure 14. Again, the transitional dynamics sharply differ from the baseline scenario with a quick reversion to the mean in formerly highly-polluted neighborhoods.

These simulations shed light on the crucial role of the *combination* of neighborhood effects and relocation rigidities in the dynamics of segregation. The Clean Air Acts of 1952 and 1968 penalized the emissions of grit, dust and “dark smoke” in urban centers, and succeeded in reducing emissions from burning coal. However, past pollution exposure leaves a long-lasting shadow. The persistence results from the inherited neighborhood composition, neighborhood effects and the subsequent persistence of neighborhood sorting.

5 Conclusion

This paper documents and presents a plausible explanation for what was, until this paper, an anecdotal observation that the east sides of formerly-industrial cities in the western hemisphere tend to be poorer than the west sides. With rising coal use in the heyday of the industrialization, pollution became a major environmental disamenity in cities. An unequal distribution of pollution exposure induced a sorting process which left lower classes in polluted neighborhoods. Our empirical analysis relies on precise pollution estimates and identifies neighborhood sorting at a highly local level: the east/west gradient reflects a drift in pollution at the city-level but the relationship between atmospheric pollution and neighborhood composition materializes at a much more local level.

We first use data from the time before coal became the major energy technology in 1817 and data around the peak time of coal use in 1881 to show that rising pollution set off the process of residential sorting. Next, we look at the long-run consequences of this initial sorting and find that neighborhood segregation is surprisingly persistent. Finding these highly persistent effects is remarkable since industrial pollution slowed down during the twentieth century and mostly stopped in the late 1960s with the introduction of a second, stricter Clean Air Act. There exists no correlation between past industrial pollution and the relatively mild contemporary pollution in England, suggesting that other forces have sustained neighborhood segregation over time. We use a quantitative model with relocation rigidities and neighborhood effects to estimate a dynamic demand equation for neighborhoods. Our structural estimates imply non-linear transitional dynamics which relate to the literature on tipping dynamics ([Card et al., 2008](#)).

Our findings hold at least two important implications. First, the success of urban policies to revitalize deprived areas may depend on the initial level of deprivation. As outlined by our findings, very deprived neighborhoods may need a larger push to attract richer residents. This observation leads to a second implication for countries like China where pollution currently presents a major challenge. Besides the well documented short-run effects of pollution exposure on health, there are long-run consequences of an uneven pollution exposure across space: pollution induces spatial inequalities that outlive de-industrialization.

References

- Anderson, Michael L.**, “As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality,” *Journal of the European Economic Association*, 2019, *forthcoming*.
- Banzhaf, H. Spencer and Randall P. Walsh**, “Do People Vote with Their Feet? An Empirical Test of Tiebout,” *American Economic Review*, June 2008, 98 (3), 843–63.
- Barreca, Alan, Karen Clay, and Joel Tarr**, “Coal, Smoke, and Death: Bituminous Coal and American Home Heating,” Working Paper 20648, National Bureau of Economic Research February 2014.
- Baum-Snow, Nathaniel and Daniel A Hartley**, “Accounting for central neighborhood change, 1980–2010,” Technical Report 2016.
- Bayer, Patrick, Robert McMillan, Alvin Murphy, and Christopher Timmins**, “A dynamic model of demand for houses and neighborhoods,” *Econometrica*, 2016, 84 (3), 893–942.
- Beach, Brian and W. Walker Hanlon**, “Coal Smoke and Mortality in an Early Industrial Economy,” *Economic Journal*, 2018, 128, 2652–2675.
- Bleakley, Hoyt and Jeffrey Lin**, “Portage and Path Dependence,” *The Quarterly Journal of Economics*, 2012.
- Brooks, Leah and Byron Lutz**, “Vestiges of Transit: Urban Persistence at a Microscale,” Technical Report 3 2019.
- Brueckner, Jan K and Stuart S Rosenthal**, “Gentrification and neighborhood housing cycles: will America’s future downtowns be rich?,” *The Review of Economics and Statistics*, 2009, 91 (4), 725–743.
- Card, David, Alexandre Mas, and Jesse Rothstein**, “Tipping and the Dynamics of Segregation,” *The Quarterly Journal of Economics*, 2008, 123 (1), 177–218.
- Chay, Kenneth Y. and Michael Greenstone**, “Does Air Quality Matter? Evidence from the Housing Market,” *Journal of Political Economy*, 2005, 113 (2), 376–424.
- Chen, Shuai, Paulina Oliva, and Peng Zhang**, “The Effect of Air Pollution on Migration: Evidence from China,” Working Paper 24036, National Bureau of Economic Research November 2017.
- Clark, G. Kitson**, *The Making of Victorian England*, Harvard University Press, 1962.
- Clay, Karen, Joshua Lewis, and Edson Severnini**, “Canary in a Coal Mine: Infant Mortality, Property Values, and Tradeoffs Associated with Mid-20th Century Air Pollution,” Technical Report 22155 April 2016.


- Cook, L. M.**, “The Rise and Fall of the *Carbonaria* Form of the Peppered Moth,” *The Quarterly Review of Biology*, 2003, 78 (4), 399–417.
- Cook, Laurence M.**, “Records of industrial melanism in British moths,” *Biological Journal of the Linnean Society*, 2018, 125, 862–866.
- Couture, Victor and Jessie Handbury**, “Urban Revival in America, 2000 to 2010,” Working Paper, National Bureau of Economic Research 2019.
- Crafts, N. F. R. and C. K. Harley**, “Output Growth and the British Industrial Revolution: A Restatement of the Crafts-Harley View,” *The Economic History Review*, 1992, 45 (4), 703–30.
- Crafts, Nicholas**, “Economic growth during the long twentieth century. Chapter 2 in Floud, R., Humphries, J. and Johnson, P.,” 2014. Cambridge University Press.
- Davis, Donald R and David E Weinstein**, “Bones, bombs, and break points: the geography of economic activity,” *American Economic Review*, 2002, 92 (5), 1269–1289.
- Depro, Brooks, Christopher Timmins, and Maggie O’Neil**, “White Flight and Coming to the Nuisance: Can Residential Mobility Explain Environmental Injustice?,” *Journal of the Association of Environmental and Resource Economists*, 2015, 2 (3), 439–468.
- Deryugina, Tatyana, Garth Heutel, Nolan H. Miller, David Molitor, and Julian Reif**, “The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction,” Technical Report 2019.
- Durlauf, Steven N.**, “A Theory of Persistent Income Inequality,” *Journal of Economic Growth*, 1996, 1, 75–93.
- , “Neighborhood Effects,” in J. Vernon Henderson and J.F. Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4, Elsevier, 2004, pp. 2173–2242.
- Freeman, Richard, Wenquan Liang, Ran Song, and Christopher Timmins**, “Willingness to Pay for Clean Air in China,” Technical Report March 2019.
- Graff Zivin, Joshua and Matthew Neidell**, “Environment, Health, and Human Capital,” *Journal of Economic Literature*, 2013, 51 (3), 689–730.
- Guerrieri, Veronica, Daniel Hartley, and Erik Hurst**, “Endogenous Gentrification and Housing Price Dynamics,” *Journal of Public Economics*, 2013, 100, 45–60.
- Hanlon, W. Walker**, “Coal Smoke and the Costs of the Industrial Revolution,” *Economic Journal*, 2019, *forthcoming*.
- Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm**, “The Making of the Modern Metropolis: Evidence from London,” *NBER Working Paper*, 2018, (25047).

- Herrnstadt, Evan, Anthony Heyes, Erich Muehlegger, and Soodeh Saberian**, “Air Pollution as a Cause of Violent Crime: Evidence from Los Angeles and Chicago,” Working Paper 2019.
- Heyes, Anthony and Soodeh Saberian**, “Air Pollution Causes Violent Crime,” Working Paper 2015.
- Hornbeck, Richard and Daniel Keniston**, “Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872,” *American Economic Review*, 2017, *107*, 1365–1398.
- Kanefsky, John and John Robey**, “Steam Engines in 18th-Century Britain: A Quantitative Assessment,” *Technology and Culture*, 1980, *21* (2), 161–186.
- Kellett, John R.**, *The Impact of Railways on Victorian Cities*, Routledge & K. Paul, 1969.
- Kuminoff, Nicolai V., V. Kerry Smith, and Christopher Timmins**, “The new economics of equilibrium sorting and policy evaluation using housing markets,” *Journal of Economic Literature*, 2013, *51* (4), 1007–1062.
- Lamb, HH**, “British Isles Weather types and a register of daily sequence of circulation patterns. Geophysical Memoir 116,” 1972.
- Lavy, Victor, Avraham Ebenstein, and Sefi Roth**, “The Impact of Short Term Exposure to Ambient Air Pollution on Cognitive Performance and Human Capital Formation,” Working Paper 20648, National Bureau of Economic Research October 2014.
- Lawton, Richard**, “An Age of Great Cities,” *The Town Planning Review*, 1972, *43* (3), 199–224.
- Lee, Sanghoon and Jeffrey Lin**, “Natural Amenities, Neighborhood Dynamics, and Persistence in the Spatial Distribution of Income,” *The Review of Economic Studies*, 2018, *85*, 663–694.
- Logan, Trevon and John M. Parman**, “The National Rise in Residential Segregation,” *Journal of Economic History*, 2017, *77* (1), 127–170.
- Maw, Peter, Terry Wyke, and Alan Kidd**, “Canals, rivers, and the industrial city: Manchester’s industrial waterfront, 1790–1850,” *The Economic History Review*, 2012, *65* (4), 1495–1523.
- Mosley, Stephen**, *The chimney of the world: a history of smoke pollution in Victorian and Edwardian Manchester*, Routledge, 2013.
- Musson, A. E.**, “Industrial Motive Power in the United Kingdom, 1800–70,” *The Economic History Review*, 1976, *29* (3), 415–39.

- Negro, Marco Del, Domenico Giannone, Marc P Giannoni, and Andrea Tambalotti**, “Global trends in interest rates,” *Journal of International Economics*, 2019, 118, 248–262.
- Nuvolari, Alessandro, Bart Verspagen, and Nick von Tunzelmann**, “The early diffusion of the steam engine in Britain, 1700–1800: a reappraisal,” *Cliometrica*, 2011, 5 (3), 291–321.
- Ravenstein, E. G.**, “The laws of migration,” *Journal of the Statistical Society of London*, 1885, 48, 167–235.
- Redding, Stephen J. and Daniel M. Sturm**, “Estimating Neighborhood Effects: Evidence from War-time Destruction in London,” Technical Report 2016.
- Rosenthal, Stuart S.**, “Old homes, externalities, and poor neighborhoods. A model of urban decline and renewal,” *Journal of Urban Economics*, 2008, 63, 816–840.
- **and Stephen Ross**, “Change and Persistence in the Economic Status of Neighborhoods and Cities,” in J. Vernon Henderson Gilles Duranton and William C. Strange, eds., *Handbook of Regional and Urban Economics*, Vol. 5 of *Handbook of Regional and Urban Economics*, Elsevier, 2015, pp. 1048–1120.
- Schelling, Thomas C**, “Dynamic models of segregation,” *Journal of mathematical sociology*, 1971, 1 (2), 143–186.
- Shaw-Taylor, L. and E. A. Wrigley**, “Occupational structural and population change. Chapter 2 in Floud, R., Humphries, J. and Johnson P.,” *The Cambridge Economic History of Modern Britain: Volume 1, Industrialisation 1700-1860*, 2014. Cambridge University Press.
- Siodla, James**, “Razing San Francisco: The 1906 disaster as a natural experiment in urban redevelopment,” *Journal of Urban Economics*, 2015, 89, 48–61.
- Villarreal, Carlos**, “Where the Other Half Lives: Evidence on the Origin and Persistence of Poor Neighborhoods from New York City 1830-2012,” Technical Report January 2014.
- Warde, Paul**, “Energy Consumption in England and Wales, 1560–2004,” 2007.
- Williamson, J. G.**, *Coping with City Growth during the British Industrial Revolution*, Cambridge University Press, 1990.
- Wrigley, E. A.**, “The PST System of Classifying Occupations,” 2010.
- Zivin, Joshua Graff and Matthew Neidell**, “The Impact of Pollution on Worker Productivity,” *American Economic Review*, December 2012, 102 (7), 3652–73.

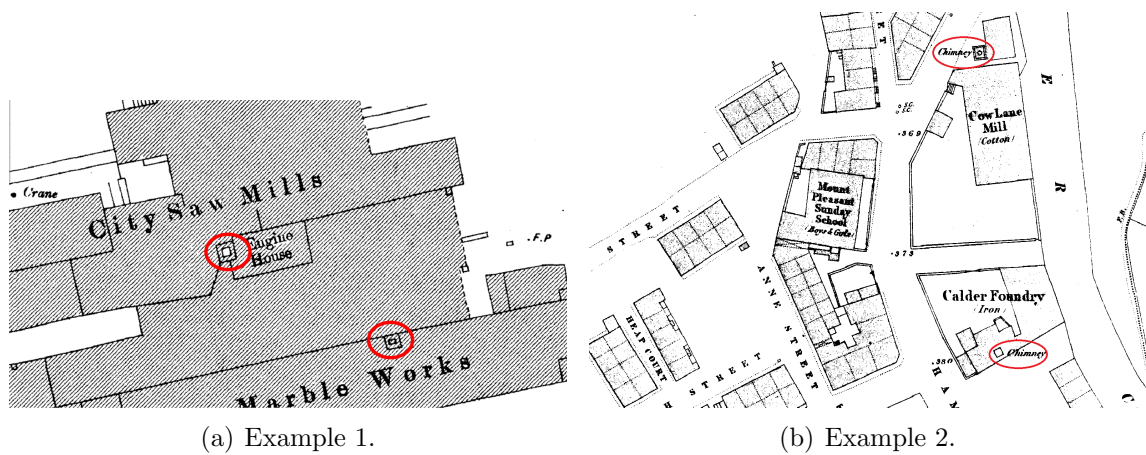
Figures and tables

Figure 1. Coal consumption and migration during the Industrial Revolution.



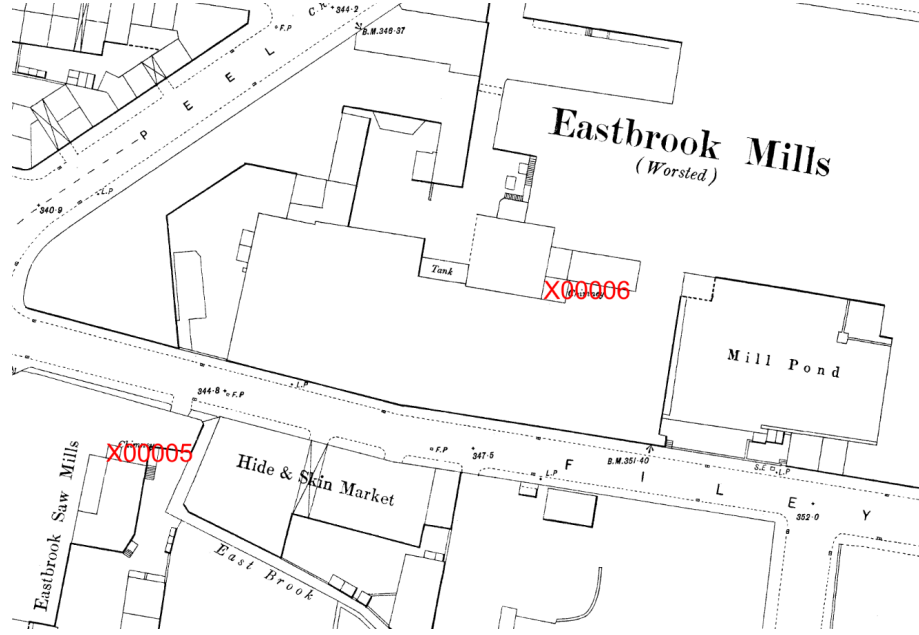
Notes: The left panel illustrates the increase and decrease in coal consumption over the period 1750–2000. The figure is based on [Warde \(2007\)](#) who reports coal consumption in petajoule. To convert numbers from petajoule to tons, we use a conversion factor of 1:34,140. The solid red lines indicate the years 1817 and 1881, while the dashed red lines mark the introduction of the 1956 and 1968 Clean Air Acts. The right panel plots the average decadal population growth rate for the period 1801–1891 in cities of our sample.

Figure 2. Ordnance Survey maps—chimney symbols.




Sources: Ordnance Survey Maps—25 inch to the mile, 1842–1952. Four different symbols for chimneys are circled.

Figure 3. Town maps—marking and identifying chimneys.



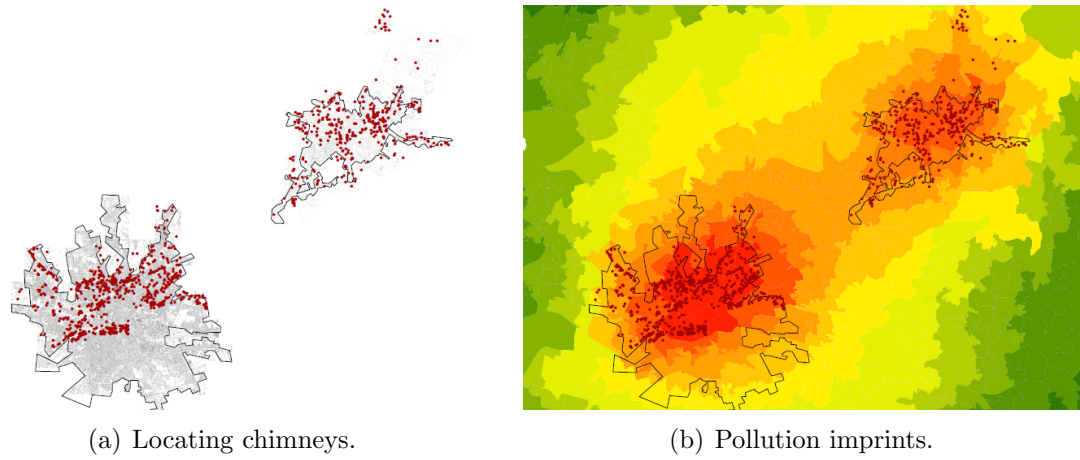
Sources: Ordnance Survey Maps—25 inch to the mile, 1842–1952. Marks X and the identifiers, e.g., 00006, are used by a recognition algorithm to locate chimneys and associate a factory.

Figure 4. Wind roses differences across two sets of meteorological conditions.




Sources: Met Office—10-year statistical meteorological data. We use 4 different sets of meteorological conditions across England and Wales: Southern England, Central England, Northern England and East Anglia.

Figure 5. Aggregating pollution sources—illustration in Manchester and Oldham.



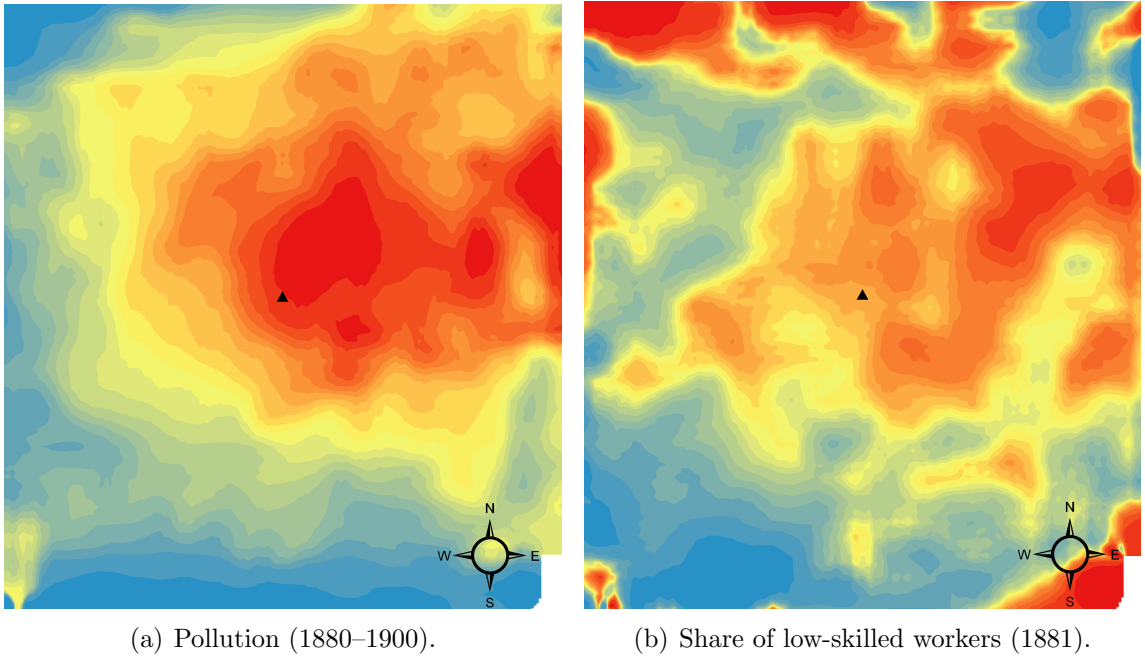
Sources: Authors' calculations using Ordnance Survey Maps—25 inch to the mile, 1842–1952 and the ADMS 5 Air Pollution Model. Chimneys are indicated with red dots.

Figure 6. Air pollution measures across neighborhoods of Manchester—external validity.




Source: First Annual Report of the Sanitary Committee on the Work of the Air Pollution Advisory Board, 1915. This Figure reports the relationship between deposits, as collected by the Air Pollution Advisory Board (1915), and our measure of SO_2 concentration in $\mu g/m^3$. The correlation between the two measures is 0.92. Deposit measure (g/m^2) is available for the following neighborhoods: Ancoats hospital (30.59), Philips Park (22.59), Whitworth Street (22.51), Queen's Park (20.18), Moss Side (18.69), Whitefield (15.53), Fallowfield (13.24), Davyhulme (12.68), Cheadle (10.63), Bowdon (6.25).

Figure 7. Pollution and shares of low-skilled workers in the average city.




Notes: The left panel of this Figure displays the average gradient of pollution in 1880–1900 across cities. The right panel displays the average gradient of neighborhood composition, as captured by the share of low-skilled workers in 1881, across cities. To construct this Figure, we define a grid of equally-spaced points within 1.5 kilometers of each city centroid (indicated with a triangle); we overlay the grids across cities and consider the unweighted average. In order to create a continuous measure, we interpolate across the grid points using a Gaussian Kernel. The 20 colors/level lines are quantiles of pollution (left panel) and share of low-skilled workers (right panel).

Figure 8. Shares of low-skilled workers and position relative to pollution sources.



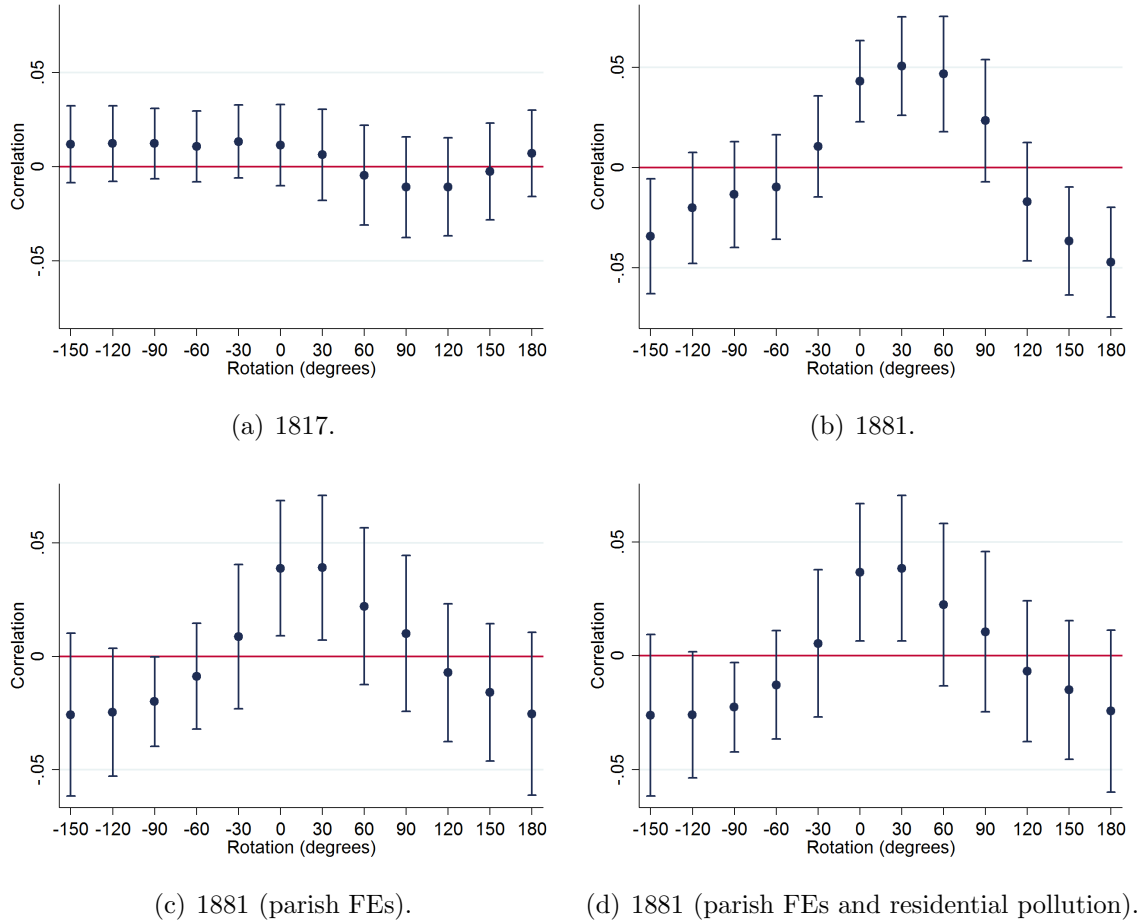
Notes: This Figure represents the relationship between the share of low-skilled workers in 1881 and the position relative to a pollution source. In this exercise, the unit of observation is a neighborhood \times chimney pair where a neighborhood is a Census cluster of households with the same geolocation in 1881 (about 150 households). The left panel displays the average share of low-skilled workers in 1881 across observed units (weighted such that all households are given the same weight). The right panel displays the residual of the share of low-skilled workers in 1881 cleaned for distance to the pollution source and distance to amenities in 1881 (distance to canals, town hall, theaters, hospitals, parks, churches, schools, universities, guild hall, mills, and elevation), and averaged over 8 main directions with respect to the pollution source. *NE* stands for north-east, indicating that the household is located toward the north-east direction, from the standpoint of the pollution source.

Figure 9. Pollution (x-axis) across neighborhoods and shares of low-skilled workers (y-axis) in 1817 and 1881.



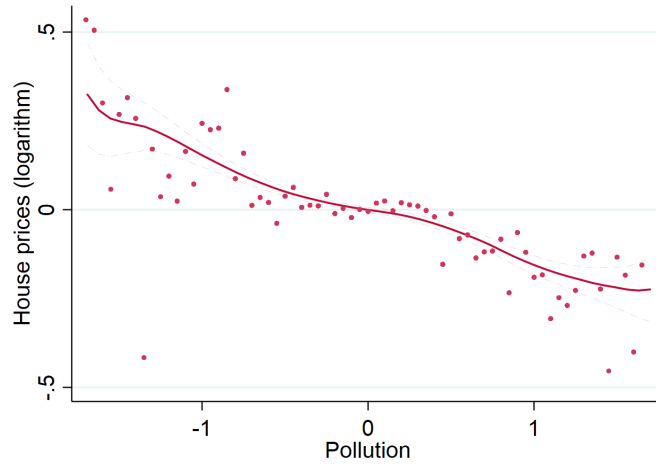
Notes: This Figure represents the relationship between the (standardized) shares of low-skilled workers in 1817 (teal, triangles) and 1881 (blue, circles) and our (standardized) measure of past pollution. We consider the residuals of all measures once cleaned by the *topography* controls, the *amenities* controls and the *lat./lon.* controls (see Table 2). We create 40 bins of neighborhoods along past pollution and the dots represent the average shares of low-skilled workers within each bin. The lines are locally weighted regressions on all observations. We restrict the sample to observations with residual pollution between -1 and 1 standard deviation(s).

Figure 10. Rotating wind patterns in 1817 and 1881.




Notes: These Figures represent the conditional correlations between the shares of low-skilled workers (in 1817 and in 1881) and counterfactual measures of past pollution rotated in steps of 30 degrees. Each dot represents the estimate in a specification including the controls reported in Table 2, Column 6, and the measure of *Static pollution* capturing proximity to the pollution source. In Panel (c), we control for parish fixed effects. In Panel (d), we control for parish fixed effects and residential pollution. Standard errors are clustered at the parish-level, and the lines represent 5% confidence intervals.

Figure 11. House transaction prices between 2000 and 2011 (y-axis) and pollution (x-axis) across neighborhoods.



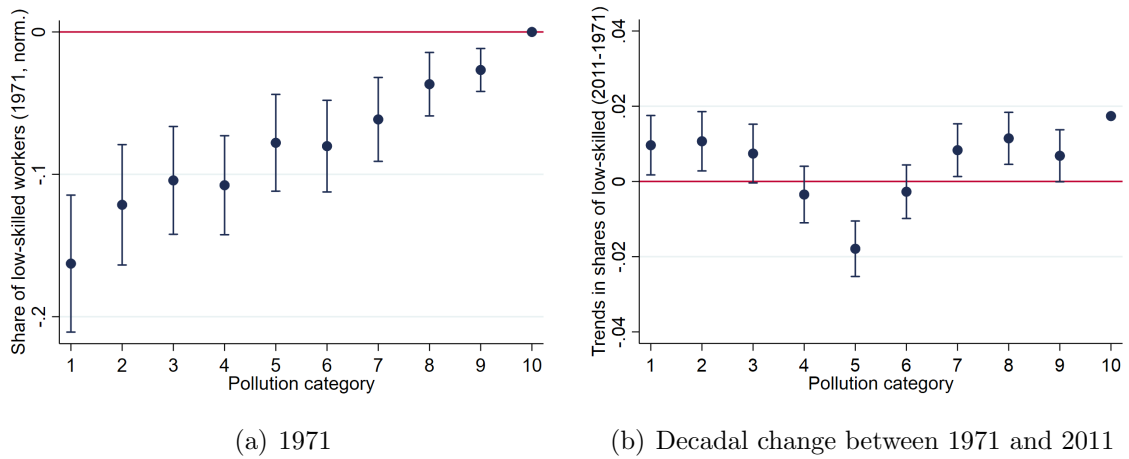
Notes: This Figure represents the relationship between the average (log) transaction prices between 2000 and 2011 and our (standardized) measure of past pollution. We consider the residuals of all measures once cleaned by city fixed effects, geographic and topographic controls. We group neighborhoods, create 100 bins of neighborhoods with similar past pollution and represent the average house prices within a pollution-bin. The lines are locally weighted regressions on all observations.

Figure 12. Pollution (x-axis) across neighborhoods and shares of low-skilled workers (y-axis) in 1817, 1881, 1971, 1991 and 2011.




Notes: This Figure represents the locally weighted regressions on all observations between the (standardized) shares of low-skilled workers and our (standardized) measure of past pollution. We consider the residuals of all measures once cleaned by city fixed effects, and *topography* and *population* controls.

Figure 13. Mean reversion between 1971 and 2011—share of low-skilled workers by pollution decile.




Notes: The left panel represents the effect of pollution on the share of low-skilled workers in 1971 in a specification with pollution category dummies defined by pollution decile (1: lowest, 10: highest pollution exposure). The effect of the last decile is normalized to 0. The right panel represents the effect of pollution on the annualized trends in low-skilled workers between 1971 and 2011. The specification is a panel regression using data in 1971, 1981, 1991, 2001, 2011 with LSOA fixed effects. The reported coefficients are extracted from the interaction of pollution category dummies and the year, and adjusted such as to represent decadal changes.


Figure 14. Simulated transitional dynamics in 1991, 2001 and 2011 using (a) the structural estimates, (b) different preferences, (c) different relocation frictions.



(a) Data (dashed) and simulation outcome (solid)



(b) No neighborhood effects



(c) Low relocation rigidities

Notes: This Figure shows the average shares of low-skilled workers as a function of past pollution, using a local polynomial smoothing. Panel (a) reports neighborhood dynamics, as estimated using the structural estimates reported in column 1 of Table 7, the observed fixed amenities—used to estimate the steady-state neighborhood composition—and the observed shares of low-skilled workers in 1971 and 1981. Dashed lines (1971, 1981, 2011 from lighter to darker blue) represent observed shares of low-skilled workers; plain lines (1991, 2001, 2011 from lighter to darker blue) represent simulated shares. Panel (b) (resp. c) displays counterfactual transitional dynamics computed with preferences that are set to be orthogonal to current neighborhood composition (resp. lower relocation rigidities, i.e., $\theta = 0.10$). We standardize the shares of low-skilled workers in panels (b) and (c) for the sake of exposition.

Table 1. Descriptive statistics and variance decomposition.

VARIABLES	Mean	Pollution		Standard deviation		
		high	low	total	between	within
<i>Air pollution</i>						
Normalized pollution	-.034	.233	-.298	.928	.596	.560
<i>Population measures</i>						
<i>1817[†]</i>						
Low-skilled workers	.785	.788	.780	.110	.090	.072
High-skilled workers	.093	.083	.102	.086	.085	.068
Farmers	.122	.128	.117	.091	.065	.047
Property tax (log)	9.93	10.04	9.82	1.25	1.04	.076
<i>1881</i>						
Low-skilled workers	.600	.627	.574	.255	.149	.235
High-skilled workers	.278	.287	.268	.241	.125	.216
Farmers	.121	.085	.157	.199	.175	.180
<i>2011</i>						
Low-skilled workers	.587	.604	.569	.173	.117	.123
High-skilled workers	.413	.395	.430	.173	.117	.123
<i>Topography controls</i>						
Maximum elevation (m)	72.8	66.2	79.3	66.0	63.2	31.8
Minimum elevation (m)	52.5	49.9	55.2	48.9	44.1	19.6
Mean elevation (m)	62.3	57.8	66.7	55.8	51.6	23.4
Distance canals (km)	6.12	5.47	6.76	14.9	19.1	1.30
<i>Amenities controls</i>						
Distance town hall (km)	4.64	4.10	5.18	5.35	4.72	1.27
Distance parks (km)	9.57	9.37	9.77	23.9	28.8	1.15
Share LSOA within city	.296	.404	.190	.417	.245	.296
Area (square km)	.939	.599	1.27	3.99	7.99	3.49
Distance heavy (km)	2.48	1.92	3.03	6.36	8.81	.972
Distance light (km)	5.31	4.93	5.69	13.5	17.6	1.15
<i>Direction</i>						
Easting (km)	453	454	452	74.9	76.9	1.90
Northing (km)	291	291	292	123	126	1.92

Notes: All statistics are computed using the baseline sample of 5,538 LSOAs. Standard deviations are decomposed into between- and within-city standard deviations. Latitude and longitude are reported as the eastings and northings in the Ordnance Survey National Grid. The samples of high- and low-within-city pollution are defined with respect to the median city pollution. [†] Shares in 1817 are computed at the parish-level, which explains the lower variance

Table 2. Pollution and shares of low-skilled workers in 1881.

Share of low-skilled	(1)	(2)	(3)	(4)	(5)	(6)
Pollution	.0440 (.0071) [.1737]	.0428 (.0072) [.1689]	.0409 (.0068) [.1614]	.0377 (.0066) [.1491]	.0354 (.0067) [.1397]	.0337 (.0069) [.1332]
Observations	5,538	5,538	5,538	5,538	5,538	5,538
Fixed effects (city)	No	Yes	Yes	Yes	Yes	Yes
Controls (population)	No	No	Yes	Yes	Yes	Yes
Controls (topography)	No	No	No	Yes	Yes	Yes
Controls (amenities)	No	No	No	No	Yes	Yes
Controls (lat./lon.)	No	No	No	No	No	Yes

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. The set of *population* controls include the parish-level shares of farmers, managers and blue-collar workers in 1817, the logarithm of the average property tax at the parish level in 1815, and total population in 1881. The set of *topography* controls include the average, maximum and minimum elevations for the LSOA and the (inverse) distance to waterways as of 1827. The set of *amenities* controls include the (inverse) distance to the city hall, the (inverse) distance to parks, the share of LSOA within the city borders in 1880, the LSOA area, the (inverse) distance to the closest heavy industry and the (inverse) distance to the closest light industry. *Lat./lon.* are the latitude and longitude of the LSOA centroid.

Table 3. Pollution and shares of low-skilled workers in 1881—IV specification.

<i>Panel A: First stage</i>		Pollution			
	(1)	(2)	(3)	(4)	
Pollution (waterways)	.2904 (.0331)	.2020 (.0347)	.3100 (.0403)	.1834 (.0384)	
<i>Panel B: Second stage</i>		Share of low-skilled workers			
	(1)	(2)	(3)	(4)	
Pollution	.1286 (.0199) [.5076]	.0936 (.0296) [.3695]	.1143 (.0201) [.4510]	.0695 (.0358) [.2743]	
Observations	4,830	4,830	4,557	4,557	
F-statistic	77.16	33.91	59.22	22.81	
OLS coefficient	.0408	.0216	.0392	.0208	
Sample	Canal>250m	Canal>250m	Canal>500m	Canal>500m	
Fixed effects (city)	Yes	Yes	Yes	Yes	
Extended controls	No	Yes	No	Yes	

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. The top panel reports the first stage, and Kleibergen-Paap F-statistics are reported in the bottom panel. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2. The variable *Pollution (waterways)* is the first predicted pollution instrument from a uniform allocation of pollution sources along waterways (as of 1827). In columns 1 and 2 (resp. 3 and 4), we exclude LSOAs within 250 meters (resp. 500 meters) of a waterway.

Table 4. Pollution and shares of low-skilled workers in 1971–2011.

Share of low-skilled workers	1971	1981	1991	2001	2011
Pollution	.0243 (.0046) [.1914]	.0309 (.0050) [.2203]	.0388 (.0063) [.2071]	.0374 (.0063) [.2298]	.0354 (.0057) [.2028]
Observations	5,535	5,538	5,538	5,538	5,538
Fixed effects (city)	Yes	Yes	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes	Yes	Yes

Standard errors are reported between parentheses and are clustered at the parish-level (as defined in 1881). Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2.

Table 5. Pollution and shares of low-skilled workers in 1971–2011—IV specification.

Share of low-skilled workers	1971	1981	1991	2001	2011
Pollution	.0300 (.0164) [.2357]	.0406 (.0198) [.2894]	.0341 (.0257) [.1823]	.0461 (.0193) [.2838]	.0494 (.0213) [.2829]
Observations	4,829	4,830	4,830	4,830	4,830
F-statistic (first stage)	33.91	33.91	33.91	33.91	33.91
Fixed effects (city)	Yes	Yes	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes	Yes	Yes

Standard errors are reported between parentheses and are clustered at the parish-level (as defined in 1881). Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2. As in column 1 of Table 3, the instrument is the predicted pollution generated by a uniform allocation of pollution sources along waterways (as of 1827), and we exclude LSOAs within 250 meters of a waterway.

Table 6. Pollution, house prices and transactions (Nationwide, 2009–2013, and Land registry, 2000–2011).

	Nationwide		Land registry	
House prices	(1)	(2)	(3)	(4)
Pollution	-.1035 (.0168) [-.1684]	-.0852 (.0121) [-.1385]	-.1116 (.0161) [-.2030]	-.0642 (.0120) [-.1168]
Observations	5,226	5,226	5,538	5,538
Fixed effects (city)	Yes	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes	Yes
Controls (house ch.)	No	Yes	No	Yes

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. Each column is the result of a separate regression. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2. The dependent variables are the (log) average house prices (from Nationwide in columns 1 and 2, and Land registry in columns 3 and 4). The unreported effect of pollution on the number of transactions is between -0.0473 (specification similar to column 3) and -0.0878 (specification similar to column 4). In column 2, controls for house characteristics include the average shares of new houses, the average square meters, number of bedrooms and the year of construction for the Nationwide transactions. In column 4, controls for house characteristics include the average shares of detached, semi-detached, terraced houses and new houses for all transactions.

Table 7. Relative demand for neighborhoods—structural estimation.

Specification (S3)	(1)	(2)
Relocation rigidity θ	.4925 (.0115)	.4965 (.0097)
Discount factor β	.9239 (.1704)	.9029 (.1841)
Relative valuation h		
Linear, $s_{j,c,t}$	-3.070 (.6395)	-2.813 (.7263)
Interaction, $s_{j,c,t}a_{j,c}$		5.666 (3.320)
Observations	16,300	16,284
Fixed effects (neighborhood)	Yes	Yes
Fixed effects (city/wave)	Yes	Yes


Standard errors are reported between parentheses. The unit of observation is a LSOA in 1981, 1991 or 2001. The estimation is performed using a one-step Generalized Method of Moments, with all endogenous variables included as instruments, except neighborhood composition and its interaction with fixed amenities which are replaced by a measure of historical pollution interacted with time, and its interaction with fixed amenities. Controls include city/wave and neighborhood fixed effects and the interactions of (i) bombing intensity during the German Blitz and (ii) the initial share of social housing in 1971, with time. The relocation rigidity corresponds to a yearly-probability to be able to relocate of 7%; the 10-year discount rate β corresponds to a yearly discount rate of 0.992 (column 1).

ONLINE APPENDIX—not for publication

A	Additional figures and tables	49
B	Appendix on the dynamic model of residential sorting	55
B.1	Equilibrium and land market clearing	55
B.2	Steady-state and stability	56
B.3	Many household types and elastic land supply	57
C	Data sources	59
D	Geolocating individuals in census data	64
E	Sensitivity analysis	68
F	Differences across neighborhoods in school supply, crime, housing quality and public amenities	76
G	Role of social housing (1971–2011)	79
H	Local factors driving the dynamics of persistence	81


A Additional figures and tables

Figure A1. Topography and historical air pollution—the example of Oldham.




Sources: ADMS 5. These maps show elevation (from green to brown) and the level lines for industrial and residential pollution in Oldham.

Figure A2. Share of darker moths (*carbonaria*) and air pollution—external validity.



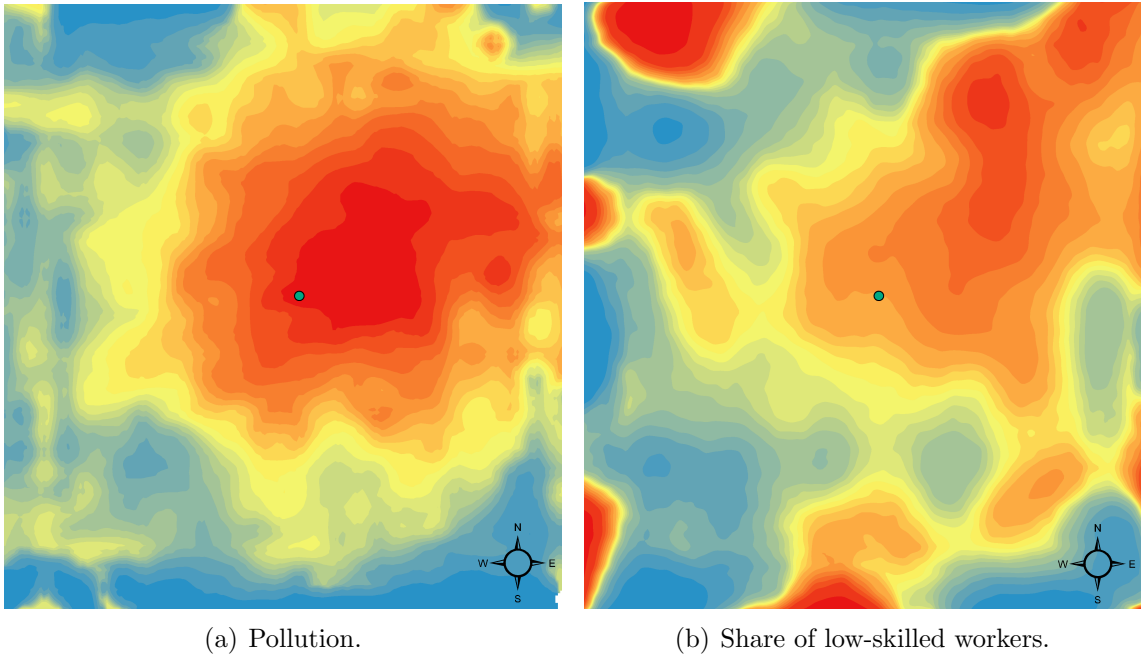
Sources: This Figure shows the relationship between the darker melanic form in species of moths and our measure of SO₂ concentration in $\mu\text{g}/\text{m}^3$. We use 54 geolocated surveys of moths, collected between 1965 and 1975 across various sites around English cities and with more than 30 observations (Cook, 2018). We construct the share of the darker form (*carbonaria*). As we condition on the species (mostly *Biston betularia* and *Odontoptera bidentata*), the y-axis represents the residual share of the darker form within species. The correlation between the two measures is 0.66, and is weighted by the number of moths identified at each survey site.

Figure A3. Cumulative of pollution in our baseline sample of 5,538 LSOAs and National Ambient Air Quality Standards (12-15 $\mu g/m^3$).



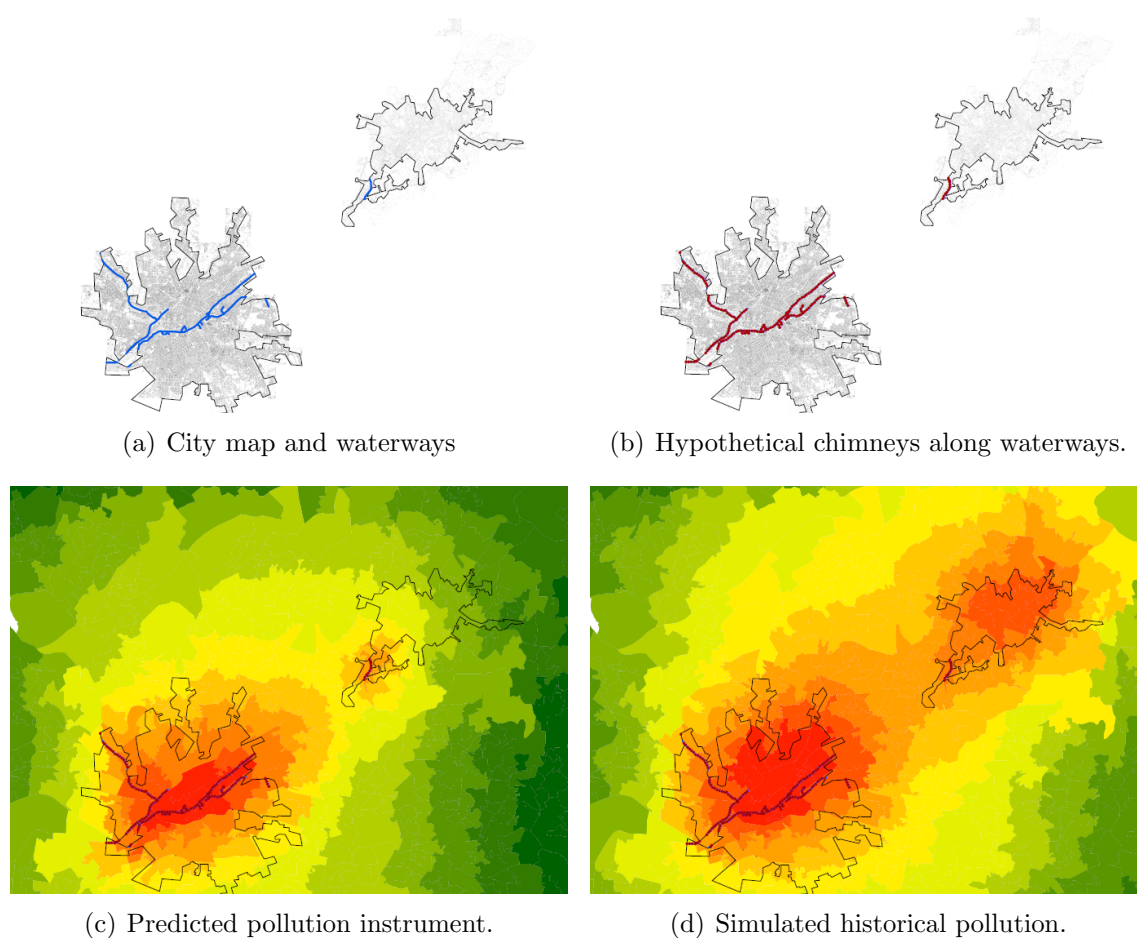
Sources: This Figure represents the cumulative distribution of SO₂ concentration as predicted by the distribution of pollution sources, the emission intensity and air pollutant dispersion. See Section 1 and Appendix Section C for additional details.

Figure A4. Pollution and shares of low-skilled workers in the average city—robustness check using town halls in order to calibrate grids across cities.



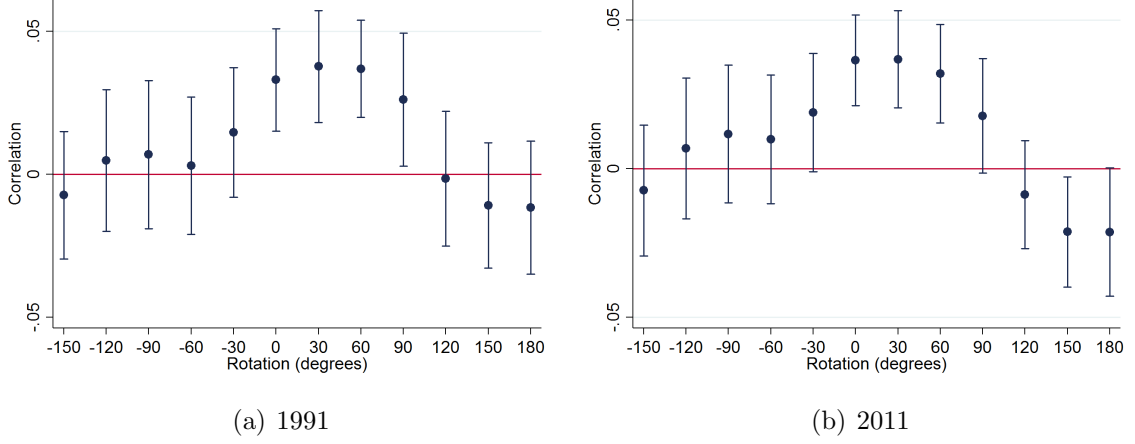
Notes: The left panel of this Figure displays the average gradient of pollution in 1880–1900 across cities. The right panel displays the average gradient of neighborhood composition—as captured by the share of low-skilled workers in 1881—across cities. To construct this Figure, we define a grid of equally-spaced points within 1.5 kilometers of the town hall of each city (indicated with a green circle); we overlay the grids across cities and consider the unweighted average. In order to create a continuous measure, we interpolate using a Gaussian Kernel. The 20 level lines are quantiles of pollution (left panel) and share of low-skilled workers (right panel).

Figure A5. The IV empirical approach with hypothetical chimneys located along old waterways—illustration in Manchester and Oldham.



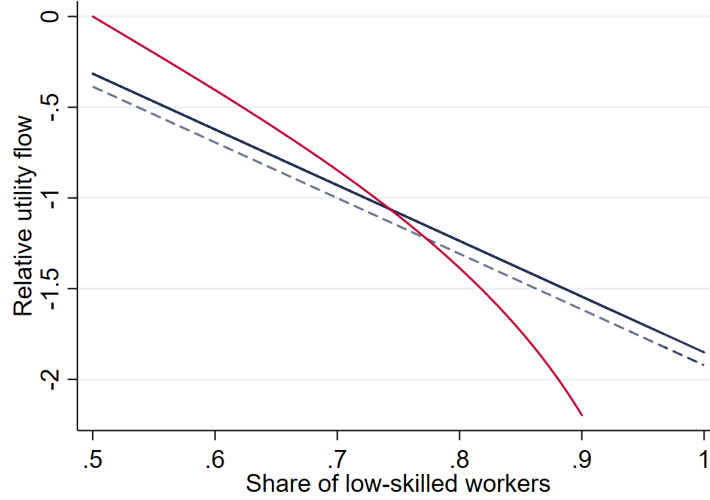
Sources: Authors' calculations using Ordnance Survey Maps—25 inch to the mile, 1842–1952 and the ADMS 5 Air Pollution Model. Chimneys are indicated with a red dot, and 1827 natural waterways with blue lines.

Figure A6. Rotating wind patterns in 1991 and 2011.



Notes: These Figures represent the conditional correlations between the shares of low-skilled workers (in 1991 and in 2011) and counterfactual measures of past pollution rotated in steps of 30 degrees. Each dot represents the estimate in a specification including the controls reported in Table 2, Column 6, and the measure of *Static pollution* capturing proximity to the pollution source (see Table A7). Standard errors are clustered at the parish-level, and the lines represent 5% confidence intervals.

Figure A7. Illustration of the steady-state equilibrium as induced by our estimates of neighborhood demand.



Notes: This Figure represents the steady-state share of low-skilled workers, $s_{j,c}$, at the equilibrium. More precisely, it represents the equation,

$$(1 - \theta\beta) \ln \left(\frac{1 - s_{j,c}}{s_{j,c}} \right) = h(a_{j,c}, s_{j,c}) + \nu_c.$$

evaluated using the estimates of Table 7, column 1. The red line represents the left-hand side of the equation while the blue line represents the right-hand side. The dashed line shows the effect of an off-equilibrium increase in the share of low-skilled workers (one standard deviation or 7 percentage points) on the relative valuation (a shift in demand) and the subsequent negative adjustment in the share of high-skilled workers (of about 3 percentage points).

Table A1. Average coal use per industry, and estimated average coal use per chimney.

Industrial categories	Average coal use C_i <i>m.tons/year</i>	Average weight E_i
Breweries	19.4	0.36
Brick factories	48.9	1.05
Chemical factories	40.1	0.84
Food processing	12.0	0.77
Foundries	43.7	1.00
Mining	28.9	3.55
Paper production	9.7	2.47
Shipbuilding	6.1	1.02
Tanneries	12.1	0.17
Textile production	10.1	0.47
Wood processing	5.4	0.10
Other manufactures	-	0.10

Source: [Hanlon \(2019\)](#) and the 1907 Census of Production. Notes: Average coal use per worker C_i is reported in tons per year and the estimated coal use per chimney E_i is normalized such as to be equal to 1 for the category “Foundries”. The measure E_i is set to the minimum average value for chimneys classified as “Other manufactures”.

Table A2. Pollution and other outcomes in 1881.

VARIABLES	Low-skilled, all	Migrants		
		All	England/Wales	Commonwealth
Pollution	.0242 (.0051) [.1192]	.0314 (.0072) [.1103]	.0282 (.0071) [.1101]	.0033 (.0019) [.0433]
Observations	5,538	4,312	4,312	4,312
Fixed effects (city)	Yes	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes	Yes

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. *Low-skilled, all* is the share of all workers between 25 and 55 years old that are employed in low-skilled occupations. Migrants are defined as individuals between 25 and 55 years old who are born in a different county. The set of extended controls include all controls of column 6 in [Table 2](#).

Table A3. Pollution and shares of low-skilled workers in 1881 and 2011—the role of covariates.

Share of low-skilled workers	1881	2011
Pollution	.0337 (.0069)	.0354 (.0057)
Employment (thousand, 1881)	-.0007 (.0005)	.0007 (.0004)
Share low-skilled (1817)	.2455 (.0888)	-.0111 (.0887)
Share farmers (1817)	-.0783 (.0954)	-.1614 (.0986)
Property tax (1815)	.0091 (.0073)	-.0111 (.0057)
Maximum elevation	-.0001 (.0004)	-.0006 (.0002)
Minimum elevation	-.0010 (.0004)	.0003 (.0004)
Average elevation	.0002 (.0007)	-.0005 (.0004)
Distance waterways (inverse)	.0705 (.0170)	.0351 (.0313)
Distance hall (inverse)	-.0083 (.0283)	.0250 (.0202)
Distance parks (inverse)	-.0306 (.0120)	-.0432 (.0104)
Share area (city)	.0202 (.0126)	-.0285 (.0135)
Area	-.0021 (.0009)	-.0002 (.0008)
Distance heavy industry (inverse)	.0017 (.0372)	-.0219 (.0631)
Distance light industry (inverse)	.3454 (.1448)	.1378 (.1163)
Longitude	.0021 (.0022)	.0064 (.0023)
Latitude	.0030 (.0023)	.0023 (.0020)
Observations	5,538	5,538

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area.

B Appendix on the dynamic model of residential sorting

In this section, we describe the detailed derivation of the dynamic demand equation, we provide a characterization of (i) the steady-state residential sorting and (ii) the stability of the dynamic system, and we extend the dynamic model to many household types and elastic land supply.

B.1 Equilibrium and land market clearing

We normalize total population to $N = 1$ and consider two types of households $o \in \{H, L\}$ in fixed aggregate proportions $N^H = \eta \in (0, 1)$ and $N^L = (1 - \eta)$.

Let $N_{j,t}^o$ denote the number of o -type households in j at t , so that $n_{j,t}^o = N_{j,t}^o/N^o$ is the share of o -type households in neighborhood j at time t . At each period t , the number of o -type households is composed of a share θ of non-movers from previous period, and a share $1 - \theta$ of households having just opted for neighborhood j . Consequently, the law of motion requires that $n_{j,t}^{*,o} = \frac{n_{j,t}^o - \theta n_{j,t-1}^o}{1 - \theta}$, and the respective demand by each type for neighborhood j is given by Equation (1),

$$\begin{cases} \ln \left(\frac{n_{j,t}^H - \theta n_{j,t-1}^H}{1 - \theta} \right) = \sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^H - \mu_t^H \\ \ln \left(\frac{n_{j,t}^L - \theta n_{j,t-1}^L}{1 - \theta} \right) = \sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{j,t+\tau}^L - \mu_t^L \end{cases}$$

where μ_t^o are type-specific aggregate present discounted values.

The land market clearing condition for each neighborhood j , $N_{j,t}^H + N_{j,t}^L = N_j$, can thus be written, using $n_{j,t}^o = N_{j,t}^o/N^o$, as,

$$\eta n_{j,t}^H + (1 - \eta) n_{j,t}^L = n_j,$$

where n_j is the fixed land supply in neighborhood j .

We use these equilibrium conditions in periods t and $t - 1$ in order to express the whole problem as a function of the share of type- L households in neighborhood j and period t , $s_{j,t} = N_{j,t}^L/N_j = (1 - \eta)n_{j,t}^L/n_j$. Note that, since $n_{j,t}^L = \frac{n_j s_{j,t}}{1 - \eta}$, we can re-write the law of motion for type- L households in terms of shares,

$$s_{j,t} = \theta s_{j,t-1} + (1 - \theta) \left(\frac{1 - \eta}{n_j} \right) n_{j,t}^{*,L},$$

from which we have an expression for $n_{j,t}^{*,L}$ in terms of $s_{j,t}$ and $s_{j,t-1}$. The land market clearing condition together with the law of motion for each type means that, in each

neighborhood j and at each period t ,

$$\eta \left[\theta n_{j,t-1}^H + (1 - \theta) n_{j,t}^{*,H} \right] + (1 - \eta) \left[\theta n_{j,t-1}^L + (1 - \theta) n_{j,t}^{*,L} \right] = n_j,$$

which, using $\eta n_{j,t-1}^H + (1 - \eta) n_{j,t-1}^L = n_j$, reduces to,

$$\eta n_{j,t}^{*,H} + (1 - \eta) n_{j,t}^{*,L} = n_j,$$

The resulting demand equations in terms of shares are,

$$\begin{cases} \ln \left(\frac{n_j}{\eta} \left(1 - \frac{s_{j,t} - \theta s_{j,t-1}}{1 - \theta} \right) \right) = \sum_{\tau=0}^{\infty} (\beta \theta)^\tau E_t u_{j,t+\tau}^H - \mu_t^H \\ \ln \left(\frac{n_j}{1 - \eta} \frac{s_{j,t} - \theta s_{j,t-1}}{1 - \theta} \right) = \sum_{\tau=0}^{\infty} (\beta \theta)^\tau E_t u_{j,t+\tau}^L - \mu_t^L \end{cases}$$

Subtracting the two demand equations yields the relative demand for neighborhood j from type- H households,

$$F \left(\frac{s_{j,t} - \theta s_{j,t-1}}{1 - \theta} \right) = \sum_{\tau=0}^{\infty} (\beta \theta)^\tau E_t (u_{j,t+\tau}^H - u_{j,t+\tau}^L) + \mu_t,$$

where $F(x) = \ln(1 - x) - \ln(x)$, and $\mu_t = \mu_t^L - \mu_t^H + \ln(\eta/(1 - \eta))$ captures the relative welfare of each type in period t . The relative demand verifies the following forward-looking equation:

$$F \left(\frac{s_{j,t} - \theta s_{j,t-1}}{1 - \theta} \right) - \mu_t = u_{j,t}^H - u_{j,t}^L + \beta \theta E_t \left[F \left(\frac{s_{j,t+1} - \theta s_{j,t}}{1 - \theta} \right) - \mu_{t+1} \right].$$

The dynamic demand Equation (2) for neighborhood j immediately derives from the previous expression.

B.2 Steady-state and stability

We now establish the properties of the recursive process described by Equation (2). First, the system admits at least one stationary point which verifies, for each neighborhood j :

$$(1 - \theta \beta) \ln \left(\frac{1 - s_j}{s_j} \right) = h(a_j, s_j) + \nu. \quad (3)$$

The left-hand side of the equation describes a strictly decreasing function with limits $-\infty$ and $+\infty$ in 0 and 1. The right-hand side of the equation describes an empirical relationship which needs to be estimated. One can conjecture that h is monotone and decreasing in s_j , with defined limits in 0 and 1 but, again, that needs not be

the case. For this reason, there may exist several stationary points (an uneven number). In our empirical application, the left-hand side equation will always be steeper than $h(\cdot)$: there will exist a single share s_j which verifies the steady-state equation for each neighborhood. Each of these shares is positively related to the average (relative) welfare of type- H households, as captured by $-\nu$. The equilibrium condition $\sum_j n_{j,t}^H = 1$ thus ensures that there is a unique allocation $\{s_j\}_{j=1,\dots,J}$, or equivalently $\{n_{j,t}^H\}_{j=1,\dots,J}$, verifying the system of Equations (3).

Second, Equation (2), for each neighborhood j , describes a non-linear recursive sequence of order 2. We linearize this equation around the steady-state in order to study its global asymptotic properties:

$$E_t[s_{j,t+1} - s_j] + b(s_{j,t} - s_j) + c(s_{j,t-1} - s_j) = 0,$$

where $b = -\frac{1}{\theta\beta} \left[1 - \beta\theta^2 + \frac{(1-\theta)h'(s_j)s_j(1-s_j)}{(1-\theta\beta)} \right]$ and $c = 1/\beta$. The characteristic polynomial associated with the sequence is $X^2 + bX + c$ which admits two real roots, (r_1, r_2) ,

$$\begin{cases} r_1 = \frac{-b - \sqrt{b^2 - 4c}}{2} \\ r_2 = \frac{-b + \sqrt{b^2 - 4c}}{2} \end{cases}$$

The (deterministic) solutions to the previous sequence are of the general form $k_1(r_1)^t + k_2(r_2)^t$. Under some estimates of $(\beta, \theta, h'(s_j))$, the roots will be ordered as follows, $0 < r_1 < 1 < r_2$. In such case, the equivalent of a “no-Ponzi condition”, anchoring future expectations in the absence of shocks, imposes that $k_2 = 0$ such that we have:

$$s_{j,t+1} - s_j = r_1(s_{j,t} - s_j) + \chi_{t+1},$$

with $E_t[\chi_{t+1}] = 0$. The initial share of households, $s_{j,0}$, would then be sufficient to fully characterize the expected dynamics of the system.

B.3 Many household types and elastic land supply

Consider I types of households populating a city, let μ_i denote the total share of households i within the city. The distribution of the idiosyncratic preference shocks implies that the share of type- i households opting for neighborhood j verifies,

$$x_{i,j,t}^* = \frac{e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{i,j,t+\tau}}}{\sum_j e^{\sum_{\tau=0}^{\infty} (\beta\theta)^\tau E_t u_{i,j,t+\tau}}} = \frac{e^{u_{i,j,t} - p_{j,t} + \beta\theta EU_{i,j,t+1} + \beta(1-\theta)V_{i,t+1}}}{\sum_j e^{u_{i,j,t} - p_{j,t} + \beta\theta EU_{i,j,t+1} + \beta(1-\theta)V_{i,t+1}}}.$$

Consequently,

$$\ln(x_{i,j,t}^*) = u_{i,j,t} - p_{j,t} + \beta\theta EU_{i,j,t+1} + \mu_{i,t},$$

where $\mu_{i,t}$, capturing $V_{i,t+1}$ and the denominator, is independent from neighborhood j and $p_{j,t}$, the rental price in neighborhood j , is independent of household type i . Letting $x_{i,j,t}$ denote the share of type- i households living in neighborhood j , we have that:

$$\ln \left(\frac{x_{i,j,t} - x_{i,j,t-1}}{1 - \theta} \right) - \beta \theta E_t \ln \left(\frac{x_{i,j,t+1} - x_{i,j,t}}{1 - \theta} \right) = u_{i,j,t} - p_{j,t} + \mu_{i,t} - E_t \mu_{i,t+1}. \quad (4)$$

This set of equations characterizes the demand for neighborhood j at period t .

We assume that land supply is elastic and verifies,

$$n_{j,t} = n_j p_{j,t}^\varepsilon,$$

with $\varepsilon > 0$. Land market clearing imposes that $\sum_i \mu_i x_{i,j,t} = n_{j,t}$ such that,

$$\ln \left(\sum_i \mu_i x_{i,j,t} \right) = \ln(n_j) + \varepsilon \ln(p_{j,t}). \quad (5)$$

The two sets of equations (4) (land demand) and (5) (land supply) perfectly characterize neighborhood sorting and its evolution over time. The estimation of this system however requires observing rental prices, or the housing stock, and their evolution over time.

C Data sources

In this section, we describe data construction of pollution sources and amenities in 1880–1900; census data in 1817 and 1881; bomb damage in 1940–1945; steam engines in 1700–1800 and waterways; and more recent data covering the period 1971–2016.

Pollution sources and amenities in 1880–1900 We rely on scans of map tiles of the collection “Ordnance Survey maps—25 inch to the mile (1842–1952)”. These maps are drawn differently across counties, covering different periods and different waves. For instance, Bedfordshire is covered by four waves, 1878–1883, 1899–1901, 1922–1924 and 1937–1940. All counties were supposed to be revised every twenty years on average, but rural maps were infrequently revised.

We consider the 70 largest industrial centers at the beginning of the nineteenth century, and select the nearest wave to the year 1890.³⁸ While the level of precision is unmatched (the positions of free-standing trees are reported), it comes at the expense of reproduction quality. The printing process involved zinc plates, and the resulting quality is not fitted for direct recognition processes or automatic map digitization.

In order to alleviate this issue, we design a recognition process which works as follows. In a first step, for each tile, we mark interesting landmarks with a recognizable sign (e.g., a red cross X) and an associated identifier (e.g., 00001), and we report information about the landmarks in a separate excel file (with information about the identifier, the type of landmarks and the name). The following landmarks are marked and digitized: Chimneys, parks, churches, town halls, schools and universities, public buildings.³⁹ In a second step, the mark is identified by a recognition algorithm as well as the associated identifier. The recognition algorithm

³⁸Below a list of the 70 metropolitan areas, with the (approximate) number of map tiles covering the metropolitan area in parentheses: Barrow-in-Furness (70 tiles), Bedford (25), Birkenhead (80), Birmingham (300), Blackburn (80), Bolton (60), Bradford (100), Bristol (180), Burnley (45), Burton upon Trent (35), Cardiff (115), Carlisle (20), Castleford (10), Chester (40), Coventry (45), Crewe (20), Croydon (20), Darlington (10), Derby (60), Dover (40), Gateshead (15), Gloucester (60), Grimsby (50), Halifax (40), Huddersfield (50), Ipswich (50), Keighley (15), Kidderminster (30), Kingston-upon-Hull (140), Leeds (140), Leicester (80), Lincoln (30), Liverpool (300), London (600), Luton (15), Macclesfield (30), Manchester (260), Middlesbrough (15), Newcastle-upon-Tyne (150), Newport (70), Northampton (60), Norwich (35), Nottingham (350), Oldham (180), Peterborough (50), Plymouth (80), Portsmouth (60), Preston (80), Reading (25), Rochdale (40), Rochester (15), Sheerness (15), Sheffield (250), Southampton (40), Stockport (10), Stockton-on-Tees (40), Stoke-on-Trent (75), Sunderland (40), Swansea (35), Swindon (30), Taunton (50), Tynemouth (80), Wallsend (50), Walsall (30), Warrington (30), Wigan (40), Wolverhampton (50), Worcester (60), York (70).

³⁹We complement this approach based on historical maps by using the English Heritage GIS Data provided by the Ordnance Survey, and geolocating monuments and listed buildings.

then associates geographic coordinates to this identifier. In a third step, we match geographic coordinates and the information stored in Excel separate files (type of landmarks, and name). This information is used to generate atmospheric pollution but also distance to industrial chimneys (Figure 8) and distance to amenities.

We also use these series of Ordnance Survey maps to create polygons of land use for the 70 cities and their outskirts, and we define city borders as the minimum polyline surrounding built-up areas. These borders are used to construct the control variable *Share area (city)*, and are used for sample selection in Appendix Table A9.

Census data in 1817 and 1881 In this project, we have obtained the authorization from the Cambridge Group for the History of Population and Social Structure to get access to the digitized micro-census of England & Wales in 1881 (about 26 Million individuals). The following variables are available: a parish code, an address, gender, age, marital status, birthplace, occupation, size of the household, number of relatives, inmates, offspring and servants, non-relatives and visitors. We use the parish and occupational classification developed by [Shaw-Taylor and Wrigley \(2014\)](#), which is already harmonized between 1817 and 1881.

We rely on a quasi-census of male occupations drawing upon 2 million observations. The Church of England kept records of baptisms, marriages, and burials in Parish registers, and the occupation of the father was reported for each baptism. Accordingly, the sample only features males who had a child in the covered period. More than 10,000 such Anglican baptism registers from 1813 to 1820 were digitized, and [Shaw-Taylor and Wrigley \(2014\)](#) associate a 1881 parish and an occupational code (following the PST structure, see [Wrigley, 2010](#)) to each observation.

We also use on the income tax levied between 1799 and 1816, and collect the 1815 property-tax assessment published at the parish level. In order to finance the Napoleonic Wars, 6 Schedules were developed as part of a generalized Income Tax. We use Schedule A assessment (tax on land income) as a proxy for land value (or rent per acre).

Bomb damage in 1940–1945 We use scans of the “Bomb Census survey records, HO 193(55-65), 1940–1945” provided by the National Archives. Only the following cities of our sample are covered: Barrow in Furness, Bedford, Birmingham, Bristol (Bath), Coventry, Derby, Dover, Gateshead, Grimsby, Ipswich, Jarrow (Wallsend), Lincoln, Liverpool (tracings only), London, Luton, Manchester (tracings only), Middlesbrough, Newcastle upon Tyne, Norwich, Nottingham, Oldham, Peterborough, Plymouth, Portsmouth (Gosport), Sheffield, Southampton, York. These Bomb Cen-

sus map tiles consist Ordnance Survey maps—Six-inch to the mile, drawn between 1919 and 1939, on which the *Ministry of Home Security Bomb Census Organisation* recorded damage sustained during bombing raids. The positions of bombs are marked by a red dot (see Appendix Figure A12).

We digitize bomb damage as follows. First, we superimpose these scans with original Ordnance Survey maps in order to geolocate each map tile. Second, we transform red dots into featured points, we draw a buffer of 10 meters around each point and associate them to the intersected LSOAs. Third, we define the dummy *Bombs* (used in Figure A13) as being equal to 1 if at least one bomb impact intersects with the LSOA.

Steam engines in 1700–1800 and waterways We use the directory of steam engines between 1700 and 1800, as computed by Kanefsky and Robey (1980) and revised by Nuvolari et al. (2011). We focus on 728 individual engines built in our cities of interest, of which about 400 are bought by collieries or textile mills. We use the location, the information about the owner and possible comments, and we precisely geolocate 547 of these steam engines.⁴⁰

We then map these locations, group them into small geographic clusters (“historical industrial districts”), define the centroid of these clusters and generate pollution dispersion from these centroids as if they were industrial chimneys with uniform pollutant emissions. The resulting measure is the variable *Pollution (steam engines)* of Appendix Table A10. The variable *Pollution (waterways)* of Table 3 is constructed as follows. We collect navigable waterways in 1827, select the intersection of these polylines with the 1890 city borders (as computed by land use in Ordnance Survey maps), and locate hypothetical chimneys every 150 meters along the within-city waterways. We then generate pollution dispersion from these hypothetical chimneys using the same dispersion process as in the baseline and uniform emissions.

Recent data (1971–2016) *Aggregate censuses at the LSOA level (1971–2011)*—We collect small area statistics from the 1971, 1981, 1991, 2001 and 2011 Censuses. We construct the shares of low- and high-skilled workers from the 1-digit Socioeconomic categories; the shares of first-generation migrants from the “Country of birth” variables; the shares of social housing and ownership using “Tenure and amenities”

⁴⁰The following cities of our sample are covered: Birmingham, Blackburn, Bolton, Bradford, Bristol, Burnley, Carlisle, Castleford, Darlington, Derby, Gateshead, Hull, Keighley, Kidderminster, Leeds, Leicester, Liverpool, London, Macclesfield, Manchester (Salford), Newcastle upon Tyne, Northampton, Norwich, Nottingham, Oldham, Plymouth, Portsmouth, Preston, Reading, Rochdale, Sheffield, Stockport, Stoke on Trent, Sunderland, Tynemouth, Walsall, Warrington, Wigan, Wolverhampton, York.

questions. We use the 2001 Lower Layer Super Output Areas (about 2,000 people per LSOA on average) as the main unit of analysis throughout the paper, and we sometimes use the Middle Layer Super Output Areas (about 10,000 people per MSOA on average) to clean for more granular fixed effects (Appendix Table A9) or to compute measures of neighborhood similarity (Figure A13).

House prices (1995–2015)—We rely on two datasets. First, we collect the HM Land Registry Transaction Data since 1995. This gives us an exhaustive register of all residential transactions in England and Wales. However, the data do not include other individual-transaction controls than the type of housing (detached, semi-detached, terraced, flat) and its age (new/old). Second, we use data from Nationwide, one of the largest mortgage provider in England and Wales, between 2009 and 2013. The Nationwide dataset includes a wide range of controls for property characteristics (e.g., the construction date, the number of bedrooms, bathrooms, garages, the square meters or heating facilities) but only covers 15% of Land Registry transactions. We use the average LSOA transaction prices in Table 6, and we use housing type as dependent variable in Appendix Table A11 and as controls in Table 6.

Building age—We collect measures of the residential dwelling ages, grouped into three age bands, and collected by the Consumer Data Research Centre. We use these measures as dependent variables in Appendix Table A11.

Education—We gather school outcomes for all primary schools from the Ministry of Education and generate LSOA measures of school supply (private schools, school value-added, teacher-pupil ratio, teacher salary, spending per student), school composition (disadvantaged pupils: defined as being either eligible for Free Schools Meals in the last six years; or looked after continuously for 1 day or more), or outcomes (the average GPA at Key Stage 2—primary education, pupils aged 7 to 11) for the period 2012–2013.⁴¹

Crime—We collect records of all criminal incidents in 2011 and their coordinates as reported by the police, and classify them into 4 categories: Anti-social behaviors including nuisance, vandalism, street drinking, littering, or vagrancy; Burglary; Drug-related crimes; and Violent crimes.⁴² The number of crimes per inhabitants in each category is used as dependent variable in Appendix Table A11.

Deprivation indices—The English Indices of Deprivation (2010) are provided by

⁴¹In order to collapse school-level indicators at the LSOA level, we proceed as follows. We compute the distance between every LSOA centroid and all the neighboring schools. We then aggregate all measures weighting each school by the inverse of the distance to the LSOA centroid.

⁴²See <http://data.police.uk>. Note that we treat these incidents irrespectively of the outcome (court decision), and compute the number of such incidents in 2011 per 100 inhabitants.

the Social Disadvantage Research Centre at the Department of Social Policy and Social Work at the University of Oxford. The main index is a weighted average of several sub-indices (Income, Employment, Health Deprivation and Disability, Education, Skills and Training, Barriers to Housing and Services, Crime, Living Environment Deprivation), some of which using primary data that are already directly included in our analysis (e.g., Key Stage 2 scores for the Education sub-index or crime occurrence for the Crime sub-index).⁴³ The composition and construction of sub-indices can be found in: <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2010>

Amenities—We construct contemporary amenities using the Ordnance Survey data on “Points of Interest” which contain the location of all public, education and leisure services. We then construct eight indices capturing the number of 1. Parks and recreation areas, 2. Theaters and museums, 3. Churches, 4. Hospitals, 5. National and local authorities, 6. Courts and police stations, 7. Transport infrastructure (bus and train stations), 8. Botanical gardens and zoos, per 100 inhabitants.

Current pollution—We collect SO₂ pollutant concentration for the United Kingdom in 2010, as published by the Department for Environment Food & Rural Affairs (DEFRA). We use the *Current pollution* measure in Appendix Table A8.

⁴³Two sub-indices may be interesting for our purpose, because they capture outcomes directly influenced by contemporary pollution. The Health Deprivation and Disability sub-index combines a measure of premature death with a morbidity/disability ratio, the emergency admission to hospital and the proportion of adults suffering from mood and anxiety disorders. The Living Environment Deprivation index captures housing quality but also air quality.

D Geolocating individuals in census data

This section describes the census structure, the fuzzy matching procedure, the clustering algorithm and a validation procedure.

Census structure There is a strong but imperfect relationship between census neighbors and true geographic neighbors that we clarify below. As we observe the parish, all our analysis will be for individuals of the same observed parish.

Let the *census identifier* i denote a transformation of the book/folio/line numbers in a systematic order. For each entry i , we can define a step function $f : i \mapsto f(i) \in \mathbb{N}$, monotonous in i .⁴⁴ The function f defines clusters among census entries. Let $n : i \mapsto n(i)$ denote the unobserved neighborhood for an individual entry i . We assume a monotonicity property for n reflecting that enumerators were recording households in a sequential manner: If $i < j < k$ and $n(i) = n(k)$, then $n(i) = n(j) = n(k)$. If two entries are in the same block n , all entries appearing between these entries also belong to the same block.

The monotonicity property is not fully sufficient to match households. Indeed, it does not allow us to observe the relationship between the values taken by blocks $\{n_j\}_j$ and census clusters $\{f_j\}_j$, and this is due to the fact that breaks in blocks cannot be observed. For instance, within a single parish, a list of entries can be:

id	i	folio	$f(i)$	block	$n(i)$	break
1.		f_1		n_1		
\vdots		\vdots		\vdots		
45.		f_1		n_1		
46.		f_1		n_2		B_1
\vdots		\vdots		\vdots		
78.		f_1		n_2		
79.		f_2		n_2		B_2

As can be seen in the previous example, there are two types of breaks in the data, one associated with a change in blocks B_1 that cannot be observed and one associated with a change in books B_2 which is observed. The true measure of a geographic cluster (i.e., a neighborhood) is n and census cluster f is an observed, imperfect proxy. In what follows, we will describe our strategy as if census clusters were a perfect representation of geographic clusters.

⁴⁴For instance, we can group lines of the same folio/book by groups of 10, or group all entries of the same folio together.

Fuzzy matching of addresses We clean addresses by deleting blanks, normalizing terms used to indicate types of roads (e.g., road, street, avenue, bow, park, square, cottage, villas, etc.) and separating the road denomination from the attributed name. We reduce the probability of georeferencing a census address incorrectly by limiting the pool of potential matches (the contemporary geolocated addresses, and monuments and listed buildings from the English Heritage GIS Data) to those which are located in the registered parish of the census observations.

The fuzzy matching procedure generates perfect matches for 20% of the total sample, and we match 30% of the total sample with precision 0.90 (at least 90% of the original string can be found in the matched address). The covariation among census entries in unmatched addresses is small which indicates that most of the matching error comes from idiosyncratic sources. However, there remains some covariation, e.g., when streets are not found in contemporary directories or when a very large “census household”, e.g., a jail, a boarding school or a guesthouse, has a poorly reported address.


In what follows, we only keep matches with a score higher than 0.90 and consider the other cases as being unmatched. We describe in the following section how we account for potential errors in the already-matched household addresses and how we geolocate the remaining households.

Recognizing clusters and inference We infer the geolocation of all households of the same census cluster f from the geolocation of a subsample of households (with potential measurement error). We start with the sample of well-matched households and apply the following algorithm to detect geographic clusters.

1. Geolocate all georeferenced households.
2. Divide a parish into 4 equal regions, depending on their position relative to the maximum and minimum latitudes and longitudes in the sample.
3. Select the region with the largest number of observations, and temporarily drop the other observations.
4. Go back to point 2. with the newly selected region.
5. Stop after few iterations, calculate the average location of remaining households, and attribute this georeference to *all* households in the same census cluster f . We then overlay these newly-identified blocks with our consistent geographic units and attribute a unique LSOA identifier to each observation.

A graphical illustration of this algorithm is provided in Figure A8 with 2 iterations. Two (resp. three) iterations already divide a parish into 16 (resp. 64) small regions. The advantage of this process is twofold: it infers georeferences for unmatched households and it reduces noise in georeferences among already-matched units.

Figure A8. Finding geographic clusters among georeferenced households in the same census cluster.




Measurement error and sensitivity analysis The previous methodology relies on two approximations. First, while census clusters reflect underlying geographic clusters, there remains a tension when aggregating entries together into a census cluster. Having more households per census cluster raises the probability to detect a geographic location. There may, however, exist breaks within a book: the first household may be interviewed in another neighborhood than the last household, for instance if the latter are interviewed by another surveyor. In order to alleviate this issue, we repeat our algorithm by generating different census clusters (grouping 1, 2, 3, 4, 5 pages together, drawing new breaks) and we compare the resulting LSOA identifier under the different specifications. Second, the exact number of iterations in the previous algorithm or the 0.90 precision threshold to exclude poorly-matched addresses may matter. In particular, when two clusters coexist within a same group of households, the previous algorithm will select one of the two clusters and ignore the presence of the other. In order to identify these outliers, we keep track of the number of households located in the selected cluster and we generate a dummy equal to one when this number is lower than 0.50. These cases represent less than 5% of all observations.

We provide a formal validation of the clustering procedure in Appendix Figure A9: we draw a random validation sample of perfectly-matched households in

1881, accounting for 20% of all perfectly-matched households; we run the clustering algorithm using only the geo-location of perfectly-matched households outside of the validation sample; and we measure the distance between the newly-attributed coordinates and the actual coordinates (see Appendix Figure A9 for the distribution of such distance within the validation sample). About 90% of the validation sample is located within 50 meters of the previous location and about 95% is within 200 meters. An urban LSOA is typically between 0.5 and 1 square kilometer, which implies that almost all households of the validation sample are attributed to the same LSOA in both instances.

Figure A9. Validation of the clustering algorithm.




Notes: The Figure represents the distance between the baseline coordinates attributed to households and the coordinates attributed by the validation exercise for households within the validation sample and with non-zero distance (22% of the validation sample). About 78% of households in the test sample are attributed similar coordinates as in the baseline. About 90% of households are within 50 meters of their previous location, about 95% are within 200 meters.

E Sensitivity analysis

This subsection presents robustness checks around the baseline specification(s).

Balance tests and validation of the 1817 quasi-census The quasi-census of 1817 relies on the occupation of fathers as reported in baptism records of the Church of England over 1813–1820; a large share of the population may be under-represented, for instance older individuals or migrants from Catholic countries. We provide in the left panel of Appendix Figure A10 a cross-validation of the baptism records data: there is a strong correlation between the two proxies for neighborhood affluence around 1815. Numerous other validation tests are provided in [Shaw-Taylor and Wrigley \(2014\)](#).

Figure A10. Validation of the quasi-Census (1817) using property taxes at the parish level in 1815.



(a) Correlation between the share of low-skilled workers in 1817 and property taxes in 1815. (b) Relationship between property taxes in 1815 and pollution.

Sources: The left panel represents the locally weighted regressions on all observations between the share of low-skilled workers in 1817 and the (log) property taxes at the parish level (1815). The right panel represents the relationship between the (log) property taxes at the parish level (1815) and our (standardized) measure of past pollution. We consider the residuals of all measures once cleaned by the *topography* controls, the *amenities* controls and the *lat./lon.* controls (see Table 2). We create 40 bins of neighborhoods along past pollution and the dots represent the average s(log) property taxes within each bin. The line is a locally weighted regression on all observations. We restrict the sample to observations with residual pollution between -1 and 1 standard deviation(s).

Figures 9 and 10 provided suggestive evidence that within-city spatial inequalities before the energy transition to coal are unrelated to the later pollution exposure. In order to further reduce concerns about biasing effects from unobserved pre-existing neighborhood characteristics, we describe a more formal balance test in Appendix Table A4. Panel A mirrors Table 2 and shows the (absence of) correlation between atmospheric pollution and the 1817 share of low-skilled workers at the parish level. We also exploit property tax data from 1815 to infer the average wealth at the parish

level and run a similar balance test in Panel B of Appendix Table A4 (see the right panel of Appendix Figure A10 for visual evidence).

Table A4. Pollution and shares of low-skilled workers or wealth measures before pollution—balance tests in 1817.

<i>Panel A: Share of low-skilled (1817)</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
Pollution	.0077 (.0100) [.0715]	.0115 (.0142) [.1063]	.0070 (.0146) [.0654]	.0021 (.0145) [.0201]	.0048 (.0187) [.0451]	.0061 (.0182) [.0567]
Observations	559	559	559	559	559	559
Fixed effects (city)	No	Yes	Yes	Yes	Yes	Yes
Controls (population)	No	No	Yes	Yes	Yes	Yes
Controls (topography)	No	No	No	Yes	Yes	Yes
Controls (amenities)	No	No	No	No	Yes	Yes
Controls (lat./lon.)	No	No	No	No	No	Yes
<i>Panel B: Property tax (log, 1815)</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
Pollution	.3254 (.0843) [.2623]	.4828 (.1277) [.3892]	.2205 (.0926) [.1777]	.1668 (.0952) [.1344]	-.0283 (.1079) [.0228]	-.0319 (.1107) [.0257]
Observations	532	532	532	532	532	532
Fixed effects (city)	No	Yes	Yes	Yes	Yes	Yes
Controls (population)	No	No	Yes	Yes	Yes	Yes
Controls (topography)	No	No	No	Yes	Yes	Yes
Controls (amenities)	No	No	No	No	Yes	Yes
Controls (lat./lon.)	No	No	No	No	No	Yes

Robust standard errors are reported between parentheses. Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a parish. The set of *population* controls include the total population in 1817. The set of *topography* controls include the average, maximum and minimum elevations for the parish and the (inverse) distance to waterways as of 1827. The set of *amenities* controls include the (inverse) distance to the city hall, the (inverse) distance to parks, the parish share within the city borders in 1880, the parish area, the (inverse) distance to the closest heavy industry and the (inverse) distance to the closest light industry. *Lat./lon.* are the latitude and longitude of the parish centroid.

Chimney height and other pollution imprints We provide a sensitivity analysis of our working assumption on the effective height of the smoke stack. We only vary the chimney height as the exit velocity and the exit temperature affect the same crucial model input, i.e., the effective height of the smoke column in the atmosphere. We report in Appendix Table A5 the relationship between pollution, as emitted by chimneys that are 15 meters high and 40 meters high, and the share of low-skilled workers in 1881, 1971, 1981, 1991, 2001 and 2011. The estimates are similar to the baseline estimates of Table 2 and Table 4.

We use a simplified and transparent measure of pollution exposure in Appendix Table A6. We create a grid of equally-spaced points over each city, count the number of chimneys within a distance to each grid point and in four different quadrants

Table A5. Pollution and shares of low-skilled workers—robustness to chimney height.

<i>Panel A: Chimney height of 15m</i>						
Share of low-skilled	1881	1971	1981	1991	2001	2011
Pollution	.0342 (.0071) [.1352]	.0276 (.0053) [.2169]	.0341 (.0060) [.2431]	.0442 (.0072) [.2359]	.0412 (.0074) [.2535]	.0388 (.0070) [.2222]
Observations	5,538	5,535	5,538	5,538	5,538	5,538
<i>Panel B: Chimney height of 40m</i>						
Share of low-skilled	1881	1971	1981	1991	2001	2011
Pollution	.0337 (.0073) [.1329]	.0265 (.0056) [.2081]	.0337 (.0064) [.2404]	.0431 (.0078) [.2302]	.0406 (.0077) [.2495]	.0382 (.0073) [.2188]
Observations	5,538	5,535	5,538	5,538	5,538	5,538

Standard errors are reported between parentheses and are clustered at the parish-level (as defined in 1881). Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2.

(North-East, North-West, South-West, South-East), and collapse the measures at the LSOA level. The first column of Appendix Table A6 shows the correlation between the share of low-skilled workers in 1881 and the number of chimneys within 2 kilometers of the average household in the four different quadrants. Only chimneys located North-West and South-West matter—and equally so—, coinciding with prevailing downwind directions. In columns 2 and 3, we split the sample between Northern and Southern cities. Chimneys located South-West are relatively more predictive of deprivation in Northern cities, possibly reflecting the Southern winds prevalent in North England (see Figure 4).

We report a comprehensive sensitivity analysis relying on other pollution imprints in Appendix Table A7. The large number of chimneys across the city implies that a measure like the distance to the closest chimney—used as a control in the benchmark specification—may not fully capture the proximity to industrial centers. Instead, we construct an atmospheric pollution profile from existing chimneys under static wind patterns, symmetric in all directions (*Static* pollution). As shown in column 1 of Appendix Table A7, the placebo atmospheric pollution measure does not affect our estimates.⁴⁵ In column 2, we present a placebo pollution pattern that varies the emission intensity rather than wind patterns. Specifically, we assume low emissions for chimneys in polluting industries and high emissions for chimney in non-polluting

⁴⁵As the measure of static pollution is positively correlated with the share of low-skilled workers when we do not control for actual pollution, this finding indicates that (i) there are more low-skilled workers close to factories but (ii) they are mostly located downwind of factories (both features were already apparent in the descriptive Figure 8).

Table A6. Number of neighboring chimneys and shares of low-skilled workers in 1881.

	(1)	(2)	(3)
Share of low-skilled workers			
Number of chimneys (NE)	.0001 (.0080)	.0024 (.0102)	-.0072 (.0122)
Number of chimneys (NW)	.0343 (.0062)	.0374 (.0096)	.0255 (.0076)
Number of chimneys (SW)	.0332 (.0079)	.0422 (.0099)	.0215 (.0115)
Number of chimneys (SE)	.0133 (.0080)	.0070 (.0092)	.0161 (.0131)
Observations	5,538	2,741	2,797
Sample	All	North	South
Fixed effects (city)	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes

Standard errors are reported between parentheses and are clustered at the parish-level (as defined in 1881). The variable *Number of chimneys (XX)* is defined as the average number of chimneys—within 2 kilometers and in the *XX* quadrant—of a randomly drawn household within a LSOA. The *North* and *South* samples are defined relatively to the median latitude of city centroids (around the latitude of Wolverhampton). The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2.

industries.⁴⁶ Including this measure does not affect our estimates and we observe a negative and insignificant coefficient on the placebo pollution measure, which shows that using information on the related industry is critical. In column 3, we control for residential pollution, which has some predictive power: the standardized effect of residential pollution is about three times smaller than the effect of industrial pollution.

We repeat this exercise with the share of low-skilled workers in 2011 as the dependent variable in Appendix Table A8.⁴⁷ In this table, however, we also add a measure of contemporary pollution, i.e., the average yearly NO2 concentration measured by DEFRA in 2015. We find that contemporary pollution has a relatively small impact on neighborhood composition in 2011. Historical pollution is far more predictive of current spatial inequalities than current exposure to air pollutants.

Sensitivity to controls, fixed effects, sample selection and clustering In Appendix Table A9, we condition for city-specific spatial gradients with respect to production locations, we discuss the choice of fixed effects, clusters, and the robustness of the estimates to the exclusion of some regions. In Panel A, we control

⁴⁶We rank industries by their coal use per worker (see Appendix Table A1), and attribute the smallest value to the most polluting industry, the second smallest value to the second most polluting industry etc.

⁴⁷Appendix Figure A6 replicates, for 1991 and 2011, the rotation exercise performed in Figure 10, and shows that deprived areas are still situated within a narrow cone around prevailing winds originating from past industrial centers.

Table A7. Pollution and shares of low-skilled workers in 1881—other pollution imprints and residential pollution.

Share of low-skilled workers	(1)	(2)	(3)
Pollution	.0431 (.0104) [.1701]	.0492 (.0137) [.1940]	.0438 (.0136) [.1729]
Static Pollution	-.0094 (.0092) [-.0369]	-.0071 (.0099) [-.0281]	-.0083 (.0099) [-.0328]
Placebo Industry		-.0110 (.0150) [-.0432]	-.0089 (.0149) [-.0351]
Residential Pollution			.0141 (.0047) [.0555]
Observations	5,538	5,538	5,538
Fixed effects (city)	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes

Standard errors are reported between parentheses and are clustered at the parish-level. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2. *Static Pollution* is the counterfactual pollution exposure using a constant average wind along all wind directions, and the same pollution sources and emissions as in the baseline. *Placebo Industry* is the counterfactual pollution exposure using the same dispersion as in the baseline, but changing emission intensity at origin. *Residential Pollution* is the pollution generated by residential emissions. See Section 2 for additional details about the construction of these variables.

for latitude and longitude (column 1), distance to the town hall (column 2), and distance to heavy industries and distance to light industries (column 3), each time interacted with city-fixed effects. These sets of controls condition for within-city geography, allowing for different dependence in distance to the town hall or the main industries (e.g., capturing different commuting facilities across cities). The effect of one standard deviation in pollution remains stable between 3.6 and 4.5 percentage points in the share of low-skilled workers. In Panel B—column 1, we report the results of our baseline specification with parish-fixed effects (about 540 in our sample). We further expand our set of fixed effects in column 2 to electoral wards (1,440 in our sample) and in column 3 to Medium Lower Super Output Areas (1,850 in our sample). The estimates remain unchanged, even when identification comes from a within-MSOA comparison. In Panel C, we report standard errors clustered at three different levels, electoral ward, MSOA and city. Standard errors increase by about 40% between the least and most conservative choice, and our baseline analysis clustered at the parish-level is at the center of this interval. In Panel D, we estimate the baseline specification on alternative samples. We exclude Greater London in column 1, the North-West including Manchester and Liverpool in column 2, and the North-East in column 3. The estimates fluctuate around the

Table A8. Pollution and shares of low-skilled workers in 2011—other pollution imprints, residential pollution and contemporary pollution.

Share of low-skilled workers	(1)	(2)	(3)	(4)
Pollution	.0364 (.0078) [.2085]	.0380 (.0108) [.2176]	.0413 (.0113) [.2365]	.0405 (.0113) [.2315]
Static Pollution	-.0010 (.0070) [-.0057]	-.0004 (.0073) [-.0024]	.0003 (.0076) [.0018]	.0004 (.0076) [.0021]
Placebo Industry		-.0028 (.0114) [-.0163]	-.0041 (.0117) [-.0236]	-.0032 (.0118) [-.0182]
Residential Pollution			-.0087 (.0046) [-.0498]	-.0090 (.0046) [-.0517]
Current Pollution				.0106 (.0037) [.0605]
Observations	5,538	5,538	5,538	5,538
Fixed effects (city)	Yes	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes	Yes

Standard errors are reported between parentheses and are clustered at the parish-level. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2. *Static Pollution* is the synthetic pollution exposure using a constant average wind along all wind directions, and the same pollution sources and emissions as in the baseline. *Placebo Industry* is the pollution exposure using the same dispersion as in the baseline, but changing emission intensity at origin. *Residential Pollution* is the pollution generated by residential emissions. *Current Pollution* is the average yearly NO2 concentration measured by DEFRA in 2015. See Section 2 for additional details about the construction of these variables.

baseline, but they remain large in all cases. In Panel E, we analyze the sensitivity of our results to the exclusion of suburbs and rural LSOAs. In columns 1 and 2, we exclude LSOAs outside of a range of 10 and 5 kilometers around the town hall, and we exclude LSOAs whose share of area within the 1880 city borders is equal to 0.5 in column 3. Even in the last instance with only 30% of our original observations, the estimate remains precisely estimated and slightly larger than in the baseline.

Alternative instrument The location decision of polluting industries in the early nineteenth century may be associated with future planned development of downwind neighborhoods. In this section, we suggest an alternative way to obtain exogenous variation in industry location, based on the historical settlements of industries. We isolate variation induced by the location of industrial districts *before* coal became a major energy source which would affect disproportionately downwind neighborhoods. To predict early industrial districts, we locate 543 early steam engines installed between 1700–1800, using data from [Kanefsky and Robey \(1980\)](#) and [Nuvolari et al. \(2011\)](#). We model uniform air pollutant emissions from early steam

Table A9. Pollution and shares of low-skilled workers in 1881—sensitivity to flexible geographic controls within cities, fixed effects, clustering and sample selection.

	Share of low-skilled workers (1881)		
	(1)	(2)	(3)
<i>Panel A: Flexible controls</i>			
Pollution	.0420 (.0078) [.1657]	.0453 (.0078) [.1789]	.0355 (.0079) [.1401]
Observations	5,538	5,538	5,538
Controls (\times City-FE)	Coordinates	Distance to hall	Distance to chimney
<i>Panel B: Fixed effects</i>			
Pollution	.0387 (.0091) [.1527]	.0377 (.0115) [.1490]	.0393 (.0118) [.1551]
Observations	5,538	5,538	5,538
Fixed effects	Parish	Ward	MSOA
<i>Panel C: Clusters</i>			
Pollution	.0337 (.0058) [.1332]	.0337 (.0066) [.1332]	.0337 (.0086) [.1332]
Observations	5,538	5,538	5,538
Clusters	MSOA	Ward	City
<i>Panel D: Sample selection (excluding regions)</i>			
Pollution	.0297 (.0070) [.1172]	.0642 (.0126) [.2536]	.0343 (.0069) [.1353]
Observations	3,661	4,395	5,247
Excluding...	London	NW	NE
<i>Panel E: Sample selection (excluding rural neighborhoods)</i>			
Pollution	.0339 (.0068) [.1340]	.0350 (.0075) [.1383]	.0425 (.0093) [.1677]
Observations	5,122	3,521	1,583
Sample	Hall<10km	Hall<5km	LSOA \cap city>0.5

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. All specifications include all controls of column 6 in Table 2. In Panel A, we control for latitude and longitude interacted with city-fixed effects (column 1), distance to the town hall interacted with city-fixed effects (column 2), and distance to heavy industries and distance to light industries interacted with city-fixed effects (column 3). A *Ward* is an electoral ward (election for local councils), and there are 1,300 wards in our sample. A *MSOA* (Middle Lower Super Output Area) is the second smallest unit in the census, and there are 1,800 MSOAs in our sample. *London* is Greater London and includes 33 districts in addition to the City of London. *NW* is the north-western region while *NE* is the north-eastern region. In Panel E, we exclude (1) LSOAs outside of a range of 10 and (2) LSOAs outside of a range of 5 kilometers around the townhall, (3) LSOAs whose share of area within the 1890 city borders is lower than 0.5.

engine locations and use the resulting atmospheric dispersion as an instrument for actual pollution, conditioning for distance to the nearest industrial chimneys. Appendix Table A10 report four sets of estimates, restricting the sample to urban

neighborhoods within a distance of 10,000 meters or 5,000 meters of a textile mill, and with or without extended controls in addition to city fixed-effects.⁴⁸ The IV estimates are remarkably similar to the ones using waterways as an exogenous factor for industry location (see Table 3).

Table A10. Pollution and shares of low-skilled workers in 1881 (IV specification)—sensitivity analysis with another instrumental variable.

<i>Panel A: First stage</i>		Pollution		
	(1)	(2)	(3)	(4)
Pollution (steam engines)	.2760 (.0493)	.1700 (.0325)	.3331 (.0599)	.1937 (.0400)
<i>Panel B: Second stage</i>		Share of low-skilled workers (1881)		
	(1)	(2)	(3)	(4)
Pollution	.0884 (.0252) [.3488]	.0901 (.0384) [.3558]	.0752 (.0244) [.2970]	.0801 (.0402) [.3160]
Observations	2,519	2,667	1,860	1,860
F-statistic	31.30	27.33	30.89	23.49
OLS coefficient	.0297	.0221	.0272	.0193
Sample	Mill<10km	Mill<10km	Mill<5km	Mill<5km
Fixed effects (city)	Yes	Yes	Yes	Yes
Extended controls	No	Yes	No	Yes

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. The top panel reports the first stage, and Kleibergen-Paap F-statistics are reported in the bottom panel. The unit of observation is a Lower Super Output Area. All specifications include all controls of column 6 in Table 2. The variable *Pollution (steam engines)* is the predicted pollution instrument which uses the location of steam engines between 1700–1800 as pollution sources. Steam engines proxy the center of historical industrial districts which were typically producing textiles. In columns 1 and 2 (resp. 3 and 4), we exclude LSOAs outside of a range of 10 (resp. 5) kilometers around a textile factory.

⁴⁸During this time, steam engines were predominantly used in textile production and collieries. Our data show that collieries are most frequently located in the hinterland of few cities, e.g., Newcastle upon Tyne, Gateshead, Leeds or Sheffield. Since we are interested in industrial districts *within* cities, we restrict our sample to urban neighborhoods around historical textile mills.

F Differences across neighborhoods in school supply, crime, housing quality and public amenities

The quantitative analysis developed in Section 3 is silent about the exact nature of the neighborhood effects that may operate, e.g., peer effects, inertia in the housing stock or the accumulation of durable public amenities constructed during the Industrial Revolution such as parks or public services. This section describes empirical facts about formerly polluted neighborhoods in 2011. While this analysis is not causal and cannot be used as hard evidence in favor of one particular channel of persistence, it helps understand the within-city distribution of consumption amenities and its relationship with past atmospheric pollution.

In order to understand the correlation between the within-city distribution of consumption amenities and past atmospheric pollution, we report the estimates for specification (S1) where we replace our benchmark indicator of neighborhood composition by measures of local amenities (see Appendix Table A11): (i) deprivation sub-indices (Panel A), (ii) a selected set of schooling and crime indicators (Panel B), (iii) characteristics of the housing stock (Panel C), and (iv) selected city amenities (Panel D). Formerly polluted neighborhoods are consistently ranked as more deprived areas across all sub-indices of deprivation. Note that the measures *Income*, *Employment* and *Education* are the most correlated with past pollution. These measures capture the incidence of low earnings, involuntary exclusion from the labor market and a lack of attainment and skills in the local population. By contrast, the correlation between *Housing*, measuring the limited physical and financial access to housing and local services, and past pollution is quantitatively small.

We then exploit more precise measures of schooling quality and crime incidence in Panel B. While the presence of private schools and the school value-added are negatively correlated with past pollution, the effects are quantitatively small. More generally, we verify in unreported tests that measures of school supply (e.g., teacher-pupil ratio, teacher salary, spending per student) are not strongly correlated with past pollution. Instead, measures capturing directly or indirectly school composition (disadvantaged students or scores) are markedly different in formerly-polluted neighborhoods. Along the same lines, burglary, drug-related and violent crimes, that tend to happen in poorest areas, are more frequent in these formerly-polluted neighborhoods in contrast to anti-social behaviors (see columns 5 to 8).

Panel C reports the correlation between past pollution and house age (columns 1 to 4). Formerly-polluted neighborhoods are not more likely to have houses constructed before 1970, 1940 or 1900, as confirmed by the average year of construction

for transactions recorded by Nationwide. However, the housing supply remains different in these areas: one standard deviation in past pollution is associated with a 5 p.p. higher prevalence of flats and a 2 p.p. lower prevalence of villas.⁴⁹

Finally, as shown in Panel D, formerly polluted neighborhoods have more parks, recreational areas and transport facilities, and less hospitals, botanical gardens or conference centers but all estimates are small in magnitude. The demand for high-quality amenities in good neighborhoods may be counteracted by high land prices. The presence of parks and recreation areas in formerly polluted neighborhoods may not only be due to low land prices but also to former industrial sites being destroyed and reclaimed in the second half of the twentieth century.

⁴⁹Note that controlling for rigid consumption amenities (e.g., provision of public services, Victorian housing stock or old private schools constructed during industrialization) does not affect the gradient between neighborhood composition and past pollution. The housing stock or the city structure does correlate with historical pollution but has limited predicting power on neighborhood composition.

Table A11. Past pollution, and deprivation measures, education and crime indicators, housing quality and amenities in 2011.

<i>Panel A: Deprivation indices</i>									
	Index	Income	Empl.	Educ.	Health	Housing	Crime	Environ.	
Pollution	.0627 (.0119) [.2435]	.0681 (.0135) [.2407]	.0491 (.0116) [.1783]	.0781 (.0120) [.2704]	.0341 (.0083) [.1346]	.0003 (.0050) [.0012]	.0273 (.0076) [.1133]	.0559 (.0131) [.2271]	
Observations	5,538	5,538	5,538	5,538	5,538	5,538	5,538	5,538	
<i>Panel B: Education and crime</i>									
	Private School (KS2)	Student Scores (KS2)	Disadvantaged Students (KS2)	School VA (KS2)	Anti-social Behaviors	Burglary	Drug-rel. Crimes	Violent Crimes	
Pollution	-.0014 (.0027) [-.0110]	-.0049 (.0012) [-.1040]	.0073 (.0013) [.0887]	-.0000 (.0001) [-.0007]	.0042 (.0044) [.0327]	.0314 (.0084) [.0925]	.0092 (.0021) [.1186]	.0654 (.0120) [.1854]	
Observations	5,538	5,538	5,538	5,538	5,538	5,538	5,538	5,538	
<i>Panel C: Housing quality</i>									
	Building 1900	Building 1970	Building 2000	Year of construction	Square meters	Bedrooms	Flats	Detached	
Pollution	.0011 (.0109) [.0043]	-.0227 (.0117) [-.1027]	.0008 (.0053) [.0047]	1.293 (1.537) [.0396]	-2.171 (.7340) [-.0866]	-.0337 (.0165) [-.0542]	.0714 (.0117) [.2649]	-.0253 (.0054) [-.1761]	
Observations	5,538	5,538	5,538	5,226	5,226	5,226	5,538	5,538	
<i>Panel D: Amenities</i>									
	Parks	Entert.	Church	Hospital	Public	Justice	Transport	Botanical	
Pollution	.0139 (.0117) [.0317]	-.0144 (.0103) [-.0357]	.0155 (.0080) [.0549]	-.0010 (.0018) [-.0125]	.0153 (.0108) [.0389]	.0028 (.0035) [.0242]	-.0001 (.0077) [-.0003]	-.0064 (.0033) [-.0392]	
Observations	5,538	5,538	5,538	5,538	5,538	5,538	5,538	5,538	

Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. Each cell is the result of a separate regression. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2. The deprivation measures are the ranks of an LSOA (0: least deprived, 1: most deprived) along the different composite sub-indices constructed with Census data, housing data, vacancies posted, schooling outcomes, the presence of public services etc. (see Section 1). KS2 stands for Key Stage 2 (age 7–11). Building 1900, 1940 and 1970 stand for the shares of dwellings constructed before 1900, between 1900 and 1970, and after 2000. Years of construction, square meters and number of bedrooms are constructed from the non-representative set of Nationwide transactions, while the shares of flats and detached houses are constructed using the exhaustive Land Registry of transactions. Panel D uses as dependent variables the number of (i) Parks and recreation areas, (ii) Theaters and museums, (iii) Churches, (iv) Hospitals, (v) National and local authorities, (vi) Courts and police stations, (vii) Transport infrastructure, (viii) Botanical gardens and zoos, per 100 inhabitants.

G Role of social housing (1971–2011)

One factor that has fostered residential segregation between 1971 and 2011 is the liberalization of social housing. In 1979, Thatcher offered social housing tenants the ‘Right to Buy’ their property, which endogenized the distribution of social housing from 1979 onward.⁵⁰ The liberalization removed support for low-skill occupations in otherwise desirable neighborhoods. As a result, some low-skilled workers chose to sell the now-valuable housing to high-skilled workers. We use the Census in 1971, 1981, 1991, 2001 and 2011 and extract a LSOA-specific share of households living in council housing, in owned properties, and the share of migrant households.

Table A12. Pollution and social housing/migrant shares (1971–2011).

Effect of pollution on ...	1971	1981	1991	2001	2011
Social housing	.0035 (.0138) [.0136] <i>.287</i>	.0191 (.0123) [.0653] <i>.358</i>	.0316 (.0097) [.1278] <i>.297</i>	.0260 (.0078) [.1190] <i>.260</i>	.0283 (.0073) [.1404] <i>.232</i>
Owners	-.0021 (.0077) [-.0082] <i>.429</i>	-.0086 (.0089) [-.0318] <i>.494</i>	-.0251 (.0074) [-.1035] <i>.580</i>	-.0247 (.0065) [-.1073] <i>.583</i>	-.0312 (.0065) [-.1399] <i>.535</i>
Migrants (New Commonwealth)	.0067 (.0034) [.1030] <i>.041</i>	.0147 (.0046) [.1730] <i>.060</i>	.0143 (.0046) [.1718] <i>.064</i>	.0172 (.0056) [.1812] <i>.085</i>	.0253 (.0075) [.2101] <i>.128</i>
Migrants (Other)	.0004 (.0014) [.0081] <i>.034</i>	-.0003 (.0012) [-.0054] <i>.035</i>	-.0008 (.0008) [-.0116] <i>.043</i>	-.0006 (.0011) [-.0098] <i>.053</i>	.0033 (.0014) [.0439] <i>.075</i>
Observations	5,534	5,538	5,538	5,538	5,538
Fixed effects (city)	Yes	Yes	Yes	Yes	Yes
Extended controls	Yes	Yes	Yes	Yes	Yes


Standard errors are reported between parentheses and are clustered at the parish-level. Standardized effects are reported between square brackets. The average value for the explained variable is reported in italic. Each coefficient is the estimate for pollution in a separate regression. The unit of observation is a Lower Super Output Area. The set of extended controls include all controls of column 6 in Table 2.

Appendix Table A12 and Appendix Figure A11 show that, while social housing was weakly correlated with past pollution in 1971, it became increasingly present

⁵⁰The United Kingdom initiated a program of social housing with the Housing of the Working Classes Acts (circumscribed to London in 1890 and extended to all councils in 1900). Council housing was the main supply of housing services for the working class, and it was typically managed by local councils. About 30% of urban households were living in council houses in 1971.

in formerly-polluted areas after 1979. More precisely, social housing appears to be distributed relatively uniformly across neighborhoods in 1971, but already aligns with historical pollution in 1991, reaching a steady-state afterwards. In parallel, the home-ownership rate decreases in areas that were formerly affected by coal pollution. We also report the correlation between past pollution and the share of immigrants in Appendix Table A12. The share of immigrants steadily increases in formerly-polluted areas with a sharp acceleration between the last two waves, which coincides with the rise in migration from poor countries. The conversion of social housing and the selective location of immigrants may be manifestations of residential segregation, but they may also have contributed to the persistence of neighborhood sorting.

Figure A11. Social housing (y-axis) and pollution (x-axis) across neighborhoods in 1971, 1981, 1991, 2001 and 2011.



Notes: The Figure represents the locally weighted regressions on all observations between the shares of social housing and our (standardized) measure of past pollution. We consider the residuals of all measures once cleaned by city Fixed-Effects, geographic and topographic controls.

H Local factors driving the dynamics of persistence

A possible shortcoming of the quantitative analysis is that we cannot identify the different mechanisms underlying the estimated neighborhood effects. While a proper analysis of the different channels would go beyond the scope of the present investigation, we now provide indirect evidence on the possible impact of urban renewal policies.

“The Blitz” and local bomb damage We collect scans of Bomb Census maps covering about 30 cities of our sample, geolocate bomb damage (indicated by a red dot—see Appendix Figure A12), and define a dummy *Bombs* equal to 1 if at least one bomb impact has been recorded within the LSOA between 1940 and 1945. We then analyze bomb damages as shock to neighborhoods that induced changes in the local occupational structure.

Figure A12. Bomb Census maps—an example in Birmingham.



Sources: The National Archives, Bomb Census survey records—HO 193(55-65), 1940–1945. Red dots indicate the location of bomb damages.

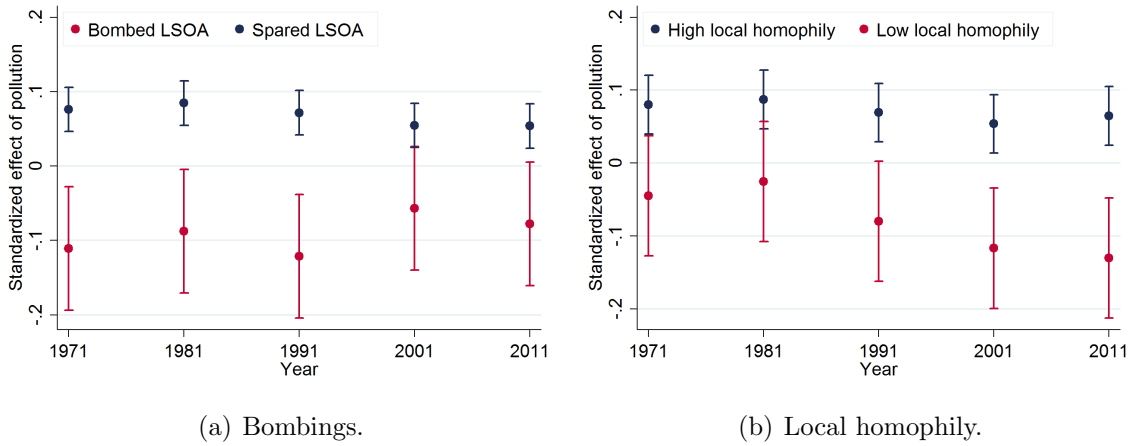
For this purpose, we run a (triple) difference-in-difference specification and identify the dynamics of persistence cleaning for location fixed effects. Letting i denote a LSOA and t a census wave ($t = 1881, \dots, 2011$), we estimate the following equation:

$$Y_{it} = \sum_{\tau=1971}^{2011} \beta_{\tau} P_i \times \mathbb{1}_{\tau=t} + \sum_{\tau=1971}^{2011} \beta_{\tau}^b P_i \times \mathbb{1}_{\tau=t} \times Bombs_i + \sum_{\tau=1971}^{2011} \gamma_{\tau} \mathbf{X}_i \times \mathbb{1}_{\tau=t} + \nu_i + \delta_t + \varepsilon_{it} \quad (\text{AS1})$$

where Y_{it} is the normalized measure of occupational structure and \mathbf{X}_i is the same set of controls as in column 6 of Table 2. In this specification, β_{τ} captures the relative long-term impact of pollution in period τ (with respect to the impact in 1881, which is the omitted category), while β_{τ}^b captures the difference between bombed and spared LSOAs in this relative long-term impact of pollution.

We report the estimates of β_{τ}^b in Panel A of Appendix Table A13. As apparent in the first column, the relative long-term impact of pollution is, in 1971, $-.19$ lower for bombed LSOAs rather than for spared locations. This gap oscillates between $-.19$ and $-.13$ from 1971 to 2011. Appendix Figure A13 illustrates the differential long-term effects of pollution, β_{τ}^b and $\beta_{\tau} + \beta_{\tau}^b$, and their evolution over time. While the impact of pollution is higher in 1971 than in 1881 for spared areas, it is markedly lower in bombed LSOAs. These findings indicate that bomb damage had a smoothing effect on the local occupational structure, and removed the long-term gradient implied by historical pollution.

Figure A13. Persistence of neighborhood sorting—role of bombings and local homophily.



Notes: The left panel represents the difference-in-difference coefficients for historical pollution in bombed LSOAs, $\beta_{\tau} + \beta_{\tau}^b$, and spared LSOAs, β_{τ} . A difference-in-difference coefficient of 0.10 implies a relative standardized effect of pollution 0.10 higher than in 1881. Bombed LSOAs are LSOAs where at least one bomb impact has been recorded between 1940 and 1945. The right panel represents the difference-in-difference coefficients for historical pollution in LSOAs with high and low homophily. A high homophily LSOA is defined as a LSOA in which normalized pollution and the average normalized pollution of its neighbors within the same MSA lie on the same side of the city average.

We consider these findings as possibly informative about the impact of urban

Table A13. Persistence of neighborhood sorting—role of bombings and local homophily.

<i>Panel A: role of bombing</i>		
	Share of low-skilled	Share of social housing
Pollution \times Bombs \times 1971	-.1870 (.0432)	
Pollution \times Bombs \times 1981	-.1719 (.0432)	-.0273 (.0255)
Pollution \times Bombs \times 1991	-.1932 (.0432)	-.0619 (.0255)
Pollution \times Bombs \times 2001	-.1112 (.0432)	-.0742 (.0255)
Pollution \times Bombs \times 2011	-.1315 (.0432)	-.0800 (.0255)
Observations	22,273	18,611
Fixed effects (LSOA)	Yes	Yes
Extended controls	Yes	Yes
<i>Panel B: role of homophily</i>		
	Share of low-skilled	Share of social housing
Pollution \times Homophily \times 1971	.1250 (.0396)	
Pollution \times Homophily \times 1981	.1126 (.0396)	-.0197 (.0237)
Pollution \times Homophily \times 1991	.1491 (.0396)	-.0311 (.0237)
Pollution \times Homophily \times 2001	.1707 (.0396)	-.0070 (.0237)
Pollution \times Homophily \times 2011	.1950 (.0396)	.0033 (.0237)
Observations	33,601	27,307
Fixed effects (LSOA)	Yes	Yes
Extended controls	Yes	Yes

Robust standard errors are reported between parentheses. All dependent variables are standardized. The unit of observation is a Lower Super Output Area \times Census wave. The set of extended controls include all controls of Table 2—column 6 interacted with wave dummies. In Panel A, *Bombs* is a dummy equal to 1 if at least one bomb impact has been recorded within the LSOA between 1940 and 1945. In Panel B, *Homophily* is a dummy equal to 1 if the LSOA normalized pollution and the average normalized pollution of its neighbors within the same MSA lie on the same side of the city average.

renewal policies and slum clearance.⁵¹ Indeed, bomb damage may be considered exogenous at the local level once controlled for distance to industries. Following

⁵¹Slum clearances in England were numerous in the early twentieth century, following the Housing Act of 1930 and the Housing Act of 1936. However, it is not possible to extract exogenous variation as they were explicitly targeting deprived but well-located neighborhoods.

destruction, the government sometimes designated these bombed locations as re-development areas with newly-built social housing, which attracted poorer workers and may have had direct long-term spillovers on neighborhood composition (Redding and Sturm, 2016).⁵² Instead, our main findings indicate that bomb damage induced redevelopment that worked against the persistence of segregation.

Local neighborhood tipping We now provide some empirical support for the theory of neighborhood tipping (Schelling, 1971). The idea is to capture how poorer LSOAs (or cells) evolve depending on the type of their immediate neighbors (or adjacent cells)—an exercise that resembles recent analyses of gentrification (Guerrieri et al., 2013; Baum-Snow and Hartley, 2016; Couture and Handbury, 2019).

We define a measure of homophily, i.e., individuals’ taste for living among their own group type, as follows. *Homophily* is a dummy equal to 1 if the LSOA normalized pollution and the average normalized pollution of its neighbors within the same MSOA lie on the same side of the city’s average pollution. High-homophily (resp. low-homophily) cells are thus surrounded by cells that are similar (resp. different) in pollution exposure.

We then run specification [AS1](#) with the dummy *Homophily* instead of the dummy *Bombs*, and we report the results in Panel B of Appendix Table [A13](#) and Appendix Figure [A13](#). The relative long-term impact of pollution is, in 1971, .12 higher in high-homophily LSOAs, and this differential steadily rises to .19 in 2011. Historically-polluted cells that are surrounded by polluted cells (about 80% of the sample) are more likely to be locked in the same equilibrium. Figure [A13](#) shows that capillarity across neighborhoods is particularly important after the Clean Air Act of 1968: while the correlation between occupational structure and pollution is already lower in low-homophily LSOAs around 1971, the gap doubles between 1971 and 2011.

The concentration of social structure has an impact on its dynamics. Air quality generated a spatially-correlated pattern in social structure within each city. Most highly polluted neighborhoods were adjacent to other polluted neighborhoods, thereby preventing occupational porosity across cells. This spatial correlation may be one component which drives the puzzling persistence of pollution effects in English modern cities. These findings may be informative about the reach of local spillovers. Urban planning and urban renewal policies may gain from targeting large areas all at once, as in Barcelona or London before their respective Olympic Games.

⁵²Note, however, that we run specification [AS1](#) for the normalized share of social housing in Panel A of Appendix Table [A13](#), and we find that social housing is less correlated with historical pollution in bombed neighborhoods.