Wages, Agglomeration and Air Pollution

Pascale Champalaune[†]

Job Market Paper

This version: 30 October 2026

Click here for the latest version

Abstract

Denser cities are more productive, but they also generate more local air pollution. This pollution lowers local amenities, which should raise real wages as compensation, but also labour productivity, which should lower them. This paper studies how these two channels shape wage and housing price differences across French cities. To do so, I develop a spatial equilibrium model in which local air pollution is generated by city population and surface area, and lowers both local amenities and labour productivity. I take it to French data over 2002-2018, using historical, geological and atmospheric instruments to estimate the urban wage and housing cost premia that separate the two channels. I find that air pollution generated by urban density lowers wages by more than housing costs, so that real wages decrease: the productivity channel dominates. As the densest cities are also the most polluted, air pollution compresses nominal wage differences across cities, but widens real wage differences. The cost of density-generated air pollution thus lies in the existing stock of density rather than in its flow. In the main counterfactual exercise, continuing current population trends while preventing cities from expanding into space, as under no net land take objectives, barely raises air pollution, and mainly raises housing prices. Reducing the elasticity of air pollution to density by a quarter would turn the resulting real wage losses into gains.

JEL Codes: R12, R13, R31, Q51, Q53, J31

Keywords: air pollution, labour productivity, amenities, density, agglomeration, wages, housing prices

[†] PSE, TSE, INRAE. Contact: pascale.champalaune@tse-fr.eu. Website: https://pchampalaune.github.io/

I am thankful to Fabian Bald, Rainald Borck, Pierre-Philippe Combes, Gabrielle Fack, Laurent Gobillon, Camille Hémet, Pamina Koenig, Katrin Millock, Eric Neumayer, Hélène Ollivier, Ariell Reshef, Sefi Roth and Gregor Singer for numerous and helpful comments and suggestions. I also thank seminar, conference, and workshop participants in Barcelona (EEA), Copenhagen (UEA), Cyprus (EAERE), London (LSE GY500 seminar), Paris (ERUDITE seminar, PSE Labour and Public Economics seminar, PSE Workshop on Cities and the Environment, AFSE, Doctorissimes), and Strasbourg (JMA) for their feedback.

This work has been funded by two French government subsidies managed by the Agence Nationale de la Recherche under the framework of the "Investissements d'avenir" programme references ANR-17-EURE-001 and ANR-18-CE22-0013-01. Access to confidential data used in this work has been made possible within a secure environment offered by CASD – Centre d'accès sécurisé aux données (Ref. 10.34724/CASD).

1 Introduction

To accommodate growing urban populations, cities are increasingly pushed to grow inward rather than outward: the European Union aims for no net land take by 2050, China requires each city to delineate an urban development boundary beyond which urban construction is restricted, and a growing number of cities now permit denser housing. Density indeed brings clear benefits: denser cities offer higher productivity and wages (Ciccone and Hall, 1996; Glaeser, 2010), and emit less CO_2 per capita (Kahn and Walsh, 2015; Lyubich, 2025).Yet denser cities are also more polluted (Borck and Schrauth, 2021; Carozzi and Roth, 2023). Unlike CO_2, a global pollutant, local air pollution harms the city's own residents: it lowers their productivity, and hence wages (Hoffmann and Rud, 2024), but, as they dislike it, it also lowers housing prices (Chay and Greenstone, 2005). Density may thus erode part of the productivity gains that make it attractive, unless lower housing prices compensate residents for the local air pollution it generates. Which of these two effects dominates is *a priori* ambiguous.

This paper quantifies the costs of the local air pollution that density generates, through its effects on wages, housing prices and real wages across cities. I develop a spatial equilibrium model in which local air pollution rises with city population but falls with city surface, and lowers both labour productivity and perceived amenities. I estimate the urban wage and housing cost premia on data for French cities over the 2002-2018 period, instrumenting density, surface area and air pollution. Matching these estimates to their model counterparts allows me to separate the productivity and amenity effects of air pollution. I then use the estimated model to simulate counterfactual scenarios that mirror current population trends and land use and air quality policies. These show that density-generated air pollution is costly as a stock: it accounts for a large share of current wage and housing price differences across cities. As a flow, however, it is not: continuing current population trends while preventing cities from expanding into space would barely worsen air quality, and reasonably weakening the link between density and air pollution would even turn the resulting real wage losses into gains.

The stock of density is costly because the productivity channel dominates: lower housing prices only partly compensate residents for local air pollution, such that real wages fall with it. Because density differs by more than an order of magnitude across cities, this drag on productivity leads to a substantial erosion of the wage gains from agglomeration: absent density-generated pollution, the urban wage premium would be significantly larger, and the cross-city variance of nominal wages about 28% higher. Current population trends, by contrast, change density by much less than it differs across cities: even when cities cannot expand into space, density rises by 13%,

but average exposure by only about 1.5%, and real wages fall mostly because housing prices rise in growing cities. The relevant policy lever is therefore the link between density and air pollution itself. Weakening it by a quarter would raise real wages by about 3% even when cities are not allowed to expand into space. While reaching the WHO air pollution guideline would require cutting this elasticity by 38%, it would increase real wages by a sizeable 7%.

To reach these conclusions, I develop a model in which workers choose which city to live in, trading off wages, housing costs and amenities (Rosen, 1979; Roback, 1982), and where to live within it. All jobs are located in the city centre, such that wages are set at the city level, while housing prices decline with distance to the centre to offset commuting costs (Alonso, 1964; Mills, 1967). Cities can therefore grow either by becoming denser or by expanding into space, with different implications for local wages and housing prices. Unlike in the canonical framework, workers are imperfectly mobile across cities: congestion costs make labour supply to each city upward-sloping, which is isomorphic to workers having idiosyncratic preferences for locations (Ahlfeldt et al., 2024). Real wage differences across cities thus reflect not only differences in amenities, but also these frictions. Firms, in turn, benefit from agglomeration economies (Duranton and Puga, 2004). The central feature of the model is that local air pollution is endogenous to city population and surface area, rather than taken as given. This builds on recent evidence that density raises local air pollution in Germany and the United States (Borck and Schrauth, 2021; Carozzi and Roth, 2023), which I confirm with original estimates for France. Local air pollution, in turn, acts as both a production and a consumption disamenity. I solve the model both when city surface is held fixed, such that population inflows are absorbed through density only, and when it adjusts endogenously. While rich, the model remains simple enough to deliver closed-form expressions for equilibrium city size, wages and housing prices.

These expressions show how density-generated air pollution affects the wage gradients that the literature estimates (Glaeser and Gottlieb, 2009). It unambiguously reduces the urban wage premium, but affects the real wage gradient in opposite directions: it lowers real wages through productivity, and raises them through compensation. How the real wage penalty to city size changes once air pollution is accounted for therefore reveals which channel dominates.

The model also disciplines the empirical strategy. Because city size, wages and housing prices are jointly determined, it implies that the wage and housing price gradients to density are identified only by instruments that shift amenities, not labour or housing construction productivity. I therefore instrument employment density with geological characteristics that determined histori-

cal settlement through soil fertility, and city surface area with pre-industrial population, since historically populated cities have on average expanded more into space. Both are plausibly unrelated to current productivity and construction costs: agriculture now accounts for about 2% of employment and value added, and population is measured in 1806, before industrialisation, conditional on long-lasting amenities such as access to the sea. I also use the two-step method of Combes et al. (2008) to purge wages of the sorting of workers across cities. Finally, because cross-city differences in air quality may correlate with other city characteristics, I instrument $PM_{2.5}$ with the number of days on which the height of the planetary boundary layer (PBL), the lowest layer of the atmosphere, falls well below the city's own long-term mean. The thinner this layer, the less room particles have to disperse upwards, and the higher the concentration of $PM_{2.5}$. Existing studies use this phenomenon to generate exogenous daily variation in local air pollution (Godzinski and Suarez Castillo, 2021; Basaglia and Sarmiento, 2025); I use it to generate yearly and cross-sectional variation in annual $PM_{2.5}$ levels.

With these instruments in hand, I proceed in three steps. First, I estimate how density and surface area affect $PM_{2.5}$, which pins down how air pollution responds to city structure in the model. In line with recent evidence for the United States (Carozzi and Roth, 2023) and Germany (Borck and Schrauth, 2021), $PM_{2.5}$ concentration increases with density, with an elasticity of about 0.12, while a larger surface area lowers it by about as much at a given population: local air pollution thus depends, to a first approximation, on density alone. Second, I estimate the urban wage and housing cost premia without controlling for air pollution, and match them to their model counterparts. As both embed the air pollution that density generates, this allows me to recover its productivity and amenity effects separately. The productivity effect dominates, such that real wages fall with air pollution: a productivity shock passes almost one-for-one into wages, whereas an amenity shock mostly capitalises into housing prices. Third, I test this prediction in reduced form, by controlling for instrumented $PM_{2.5}$. At a given level of air quality, and as unobserved heterogeneity is already controlled away, the elasticity of nominal wages to density is about 40% larger, a magnitude comparable to the role played by the spatial sorting of skilled workers in inflating the urban wage premium (Combes et al., 2008). In addition, about a third (36%) of the real wage penalty to city size operates through the air pollution that density generates when city surface is held fixed, against less than 10% when it is allowed to vary.

Beyond the reduced-form relationships, I then take the model to the data. Using the closed-form expressions for the urban wage and housing cost premia, I map the reduced-form estimates into the elasticities of labour productivity and perceived amenities to air pollution. The calibrated

model indicates that, in addition to being a consumption disamenity, which implies that some compensation for bad air quality does occur, local air pollution is an even stronger production disamenity. Real wages therefore decrease with local air pollution, because the negative productivity effect of (density-induced) air pollution is stronger in absolute terms than the compensation mechanism. This contradicts the common view that bad air quality is a pure consumption disamenity, fully compensated for through higher real wages. The dominance of the productivity channel stems only partly from the damage elasticities, as the productivity one is less than twice as large as the amenity one. It mostly stems from how each channel passes through to real wages. A loss in productivity lowers wages almost one-for-one, while housing prices fall by too little to offset it, such that real wages fall. A loss in amenities, by contrast, barely affects wages and mostly lowers housing prices, which raises real wages, but by much less: the elasticity of real wages is about 0.7 to productivity, against about 0.1 to amenities.

Then, I take the spatial equilibrium observed across French cities in 2018 as my baseline to conduct counterfactual analyses. For each exercise, I compare a fixed city surface area, which corresponds to the shorter-run response or to a binding constraint on land take, with an endogenous one, which corresponds to the longer-run response once cities can expand into space. I examine three exercises in turn.

First, I sever the link between density and air pollution and examine the effect this has on the variance of equilibrium wages and housing prices. I find that density-generated air pollution compresses nominal wages: without it, their cross-city variance would be about 28% higher, almost entirely through the productivity channel. As housing prices are much less affected, it instead widens the cross-city variance of local real wages, by 4 to 5%.

A second exercise shows that bringing average exposure within the WHO guideline of 5 $\mu g/m^3$, against about 11 $\mu g/m^3$ in 2018, would require cutting the elasticity of $PM_{2.5}$ to density by 38%, for instance through measures that reduce per-capita emissions more in denser cities, such as public transport or collective heating networks. Accounting for the resulting reallocation of workers across cities, this would raise real wages by about 7%, almost entirely through labour productivity gains. A smaller, 25% reduction would leave exposure above the guideline, but still raise real wages by about 4.5%.

I a third exercise, I project the 1999-2018 population trend forward, by backing out the corresponding perceived amenity change, and feeding it to the 2018 equilibrium. When cities cannot expand into space, as under no net land take objectives, density rises by 13.4%, but average exposure to air pollution by only about 1.5%. Real wages fall by 1.5%, mostly because housing prices

rise in growing cities, not because air quality deteriorates. Cutting the elasticity of air pollution to density by a quarter reverses this, raising real wages by 3.1% even under a fixed surface area.

All in all, these results show that accounting for the air pollution generated by density matters for both the level and the spatial distribution of wages and real wages, and that weakening this link would yield large gains, however future population is distributed across space.

These results are directly relevant to urban planning policy, in particular in a context of pressing environmental concerns. Densification is often promoted because dense cities emit less CO_2 per capita (Kahn and Walsh, 2015), but may be feared to come at the expense of worse local air quality. The results suggest that this concern applies to existing density rather than to densification. The air pollution that density generates is costly because density already differs widely across cities: accommodating current population growth within existing city limits would barely worsen air quality, and its costs would run mostly through housing prices. The relevant policy margin is therefore not how much cities densify, but how much air pollution a given density generates. Denser cities also provide local public services at a lower cost per capita, and their residents rely less on cars (Ahlfeldt and Pietrostefani, 2019), such that measures that cut emissions per resident are likely to be more cost-effective there, be it public transport investment or the replacement of individual residential heating, the largest source of $PM_{2.5}$ (Citepa, 2018), by collective heating networks. Densification and emission-reducing investments are thus complements rather than substitutes, and the latter would raise real wages even where land take is restricted. Under such restrictions, however, cleaner air makes growing cities more attractive, such that these investments should come with policies that expand housing supply within existing city limits.

This paper contributes to three strands of the literature. First, it contributes to the literature on spatial differences in wages and housing costs, by showing that local air pollution is an endogenous cost of agglomeration, which acts through labour productivity. In spatial equilibrium, cross-city differences in wages and housing costs reflect differences in productivity and amenities (Rosen, 1979; Roback, 1982; Glaeser and Gottlieb, 2009). A large literature documents the nominal wage gains from agglomeration since Ciccone and Hall (1996), and the role of the sorting of skilled workers in inflating them (Glaeser and Mare, 2001; Combes et al., 2008), including in the French context (Combes et al., 2010; Berson et al., 2024), while a more recent strand quantifies the urban costs that offset these gains, mostly through housing (Combes et al., 2019; Chen et al., 2024). I show that density also generates a cost that operates through productivity rather than through housing: the air pollution it generates erodes part of the very gains that agglomeration

produces, to an extent comparable to the bias induced by the sorting of skilled workers (Combes et al., 2008), and accounts for part of the real wage penalty to city size. Moreover, cross-city gaps in real wages are commonly interpreted as compensating differentials for differences in amenities, including climate (Albouy, 2016; Albouy et al., 2016; Diamond, 2016), and have received growing attention as a dimension of spatial inequality (Moretti, 2013; Ganong and Shoag, 2017; Diamond and Moretti, 2026). Building on recent work showing that mobility frictions matter for this interpretation (Ahlfeldt et al., 2024), I allow workers to be imperfectly mobile across cities, such that real wage gaps reflect both amenities and these frictions. I show that they also partly reflect productivity losses: because the most productive cities are also the densest and the most polluted, local air pollution compresses nominal wage differences across cities, but widens differences in real wages.

Second, this paper contributes to the literature on city structure and the environment. Denser cities emit less CO_2 per capita (Kahn and Walsh, 2015; Blaudin de Thé et al., 2021; Lyubich, 2025), and density brings a range of economic benefits (Ahlfeldt and Pietrostefani, 2019). Recent work shows that density also raises local air pollution (Borck and Schrauth, 2021; Carozzi and Roth, 2023), and discusses its implications for health. I quantify the costs of this "dirty density" beyond health, in general equilibrium, through its effects on wages, housing prices and real wages. I show that these costs lie in the existing stock of density rather than in current densification: because current population trends change density by much less than it differs across cities, accommodating them within existing city limits barely raises air pollution, and lowers real wages mostly through housing costs. The relevant policy margin is therefore how much air pollution a given density generates, rather than density itself. As total urban population is then pinned down by projected trends, the model also delivers welfare changes: weakening the link between density and air pollution by a quarter would raise welfare by 7%.

Third, this paper contributes to the literature on the economic costs of air pollution, by recovering its productivity and amenity effects jointly, in a long-run spatial equilibrium. One strand values air quality through its capitalisation into housing prices and its effect on location choices, thus treating it as a consumption amenity (Chay and Greenstone, 2005; Bayer et al., 2009). Another documents negative effects of air pollution on worker productivity and labour supply, mostly over short exposure periods and using narrow measures of output (Graff Zivin and Neidell, 2012; Chang et al., 2016, 2019; Hanna and Oliva, 2015; Hoffmann and Rud, 2024; Leroutier and Ollivier, 2025),[1] while effects on earnings have only been studied through very short-term pollution events

[1]These effects hold both for workers completing manual tasks (He et al., 2019; Adhvaryu et al., 2022) and for high-skilled workers performing cognitive ones (Künn et al., 2023; Holub and Thies, 2023; Lichter et al., 2017; Qin et al., 2022), in line with negative effects on cognitive performance (Ebenstein et al., 2016; Carneiro et al., 2021; Heissel et al.,

(Borgschulte et al., 2024).[2] I find that the labour productivity cost dominates: lower housing prices compensate residents for only part of it, such that real wages fall with air pollution. Valuations of air quality based on housing markets therefore capture only part of its cost. These effects arise in a low-pollution setting, where average exposure is close to the 2005 WHO guideline, consistent with evidence that air pollution remains harmful below legal standards (Deryugina et al., 2019; Burnett et al., 2018), and operate through the price of labour rather than through hours worked.

The remainder of the paper is organised as follows. Section 2 develops a spatial equilibrium model of the urban system with endogenous air pollution. Section 3 gives a description of all data sources. Section 4 discusses the identification strategy. Section 5 presents the reduced-form results, first on the relationship between city structure and $PM_{2.5}$ in French cities, then on the implications of air pollution for the nominal wage, housing cost and real wage differentials of agglomeration. Section 6 combines the model with the reduced-form estimates to recover the productivity and amenity effects of air pollution, and Section 7 presents the counterfactual exercises. Section 8 concludes.

2 A spatial equilibrium model with endogenous air quality

I start by motivating the main empirical analysis by setting up a model that extends the canonical framework of cross-city equilibrium of Rosen (1979) and Roback (1982). I start by relating it to the monocentric city model in the tradition of Alonso (1964) and Mills (1967). Then, I introduce several key components, starting with agglomeration economies, on the firm side, and congestion costs, or mobility frictions, on the worker side. Crucially, I then model local air pollution as being generated by local city structure, and as affecting both local perceived amenities and local productivity. I then derive closed-form equilibrium expressions for city size, wages and housing prices. The central assumption here is that equilibrium is achieved both within each city, across all possible residential locations, and across cities, such that workers are indifferent between living at x or x' within city c, and between living in city c or in city c'. Using this framework, I show theoretically how air pollution, as a product of city size, affects wage and housing cost differentials. I also discuss implications for the empirical identification of underlying parameters.

2022; Brehm et al., 2025; Krebs and Luechinger, 2024; La Nauze and Severnini, 2025).

[2] I use daily measures of PBL height, in order to build yearly-level instruments for $PM_{2.5}$. The instrument thus generates yearly variation in local air pollution. Although it does so by relying on short-term events, it does not identify a short-term (e.g., daily) effect of air pollution on wages, but a more medium term (and cross-section) effect. In that sense, this paper differs from Borgschulte et al. (2024), which looks at the effect on earnings of an additional *day* of exposure to a wildfire.

2.1 Model set-up

I focus on a circular city, which is made up of locations that are each at distance x from the central business district (CBD), where $x = 0$. Residents/workers choose to live in any location x, and since all production takes place in the CBD, they commute from x to the CBD to work. Changes in the city's surface area are therefore changes in residential extent around a fixed employment centre, whose costs fall on households through commuting rather than on firm productivity.[3]

Worker side First, the representative worker that lives in location x is assumed to have the following (normalised) Cobb-Douglas utility:

$$U(C(x), h(x)) = \frac{1}{\beta^{\beta}(1-\beta)^{1-\beta}} \Theta h(x)^{\beta} C(x)^{1-\beta} \tau(x)^{-1} \tag{1}$$

where Θ represents consumption amenities. First, I set $\Theta = \theta N^{-\delta}$, with $\delta \geq 0$ and N the entire city population, so as to model overall congestion forces. The canonical Rosen-Roback framework assumes that labour supply is perfectly elastic. But introducing $N^{-\delta}$ to the utility function implies instead that each worker inelastically supplies labour to the dedicated market located in the CBD. This is a reduced-form way of generating an upward-sloping labour supply, which is particularly important to better fit the observed spatial equilibrium.[4] Workers obtain a wage w by working at the CBD.[5]

θ captures the net benefits of living in the city, once congestion costs have been accounted for. It can be thought of as a broadly defined public good containing mostly natural amenities, such as agreeable weather. As I will discuss more thoroughly below, city-level air pollution, denoted Z, will negatively affect θ, since clean air is positively valued by workers. The rest of the utility function is made up of $h(x)$, the quantity of housing that the worker consumes at unit price $q(x)$,

[3] Liotta et al. (2022) test the relevance of the monocentric city model on a set of 192 cities across the globe, and find that its predictions capture observed within-city structure very well, in both developed and developing contexts. Harari (2020) finds that urban compactness affects household welfare through within-city commuting distances, but not firm productivity, in the Indian setting.

[4] Another way of introducing such mobility frictions is adding to the utility function an idiosyncratic amenity shock drawn from a Type-1 extreme-value distribution. In such a case, the shape parameter that controls the dispersion of the amenity draws, denoted ϵ in Ahlfeldt et al. (2015) and γ in Ahlfeldt et al. (2024), enters the labour supply equation isomorphically as in the reduced-form approach I take, with $\epsilon = 1/\delta$.

[5] I interpret w as an hourly wage and do not model the division of time between work and leisure. The model thus captures the extensive, locational margin of labour supply (that is, mobility across cities, which pins down city population N) and abstracts from the intensive margin of hours. This is a deliberate choice about the structural object: the model prices the hourly return to labour, which maps to productivity, whereas total earnings would additionally embed an hours response. Such a response is present in my setting, but it is a labour supply channel that is distinct from the price of labour, and loading it onto the structural object would misattribute an income effect response to productivity. In addition, the intensive margin is in any case compressed in my empirical setting, where statutory working time is standardised around a 35-hour reference week, limiting how much full time hours can vary across cities.

$C(x)$, a numéraire tradable good, and $\tau(x)$, a commuting disutility that the worker experiences from travelling from their residence x to their workplace, the CBD. Combining all these elements, the indirect utility function of a worker living in x writes:

$$V(x) = \theta \frac{w}{\tau(x) q(x)^{\beta} N^{\delta}} \tag{2}$$

Assuming that the economy is at spatial equilibrium, workers are indifferent between locating at any x within the city, such that $V(x) = \bar{v}$. In particular, rearranging (2) allows to write the following bid-rent equation:

$$q(x) = \left(\frac{\theta w}{\bar{v} \tau(x) N^{\delta}} \right)^{\frac{1}{\beta}} \tag{3}$$

At this equilibrium, workers are indifferent between any two locations, as distance to the CBD increases, commuting costs increase, and housing prices adjust downwards. Valuing equation (3) at the CBD and relying on utility equalisation generates the following expression for inverse labour supply:

$$w = \theta q(x)^{\beta} \tau(0) N^{\delta} \tag{4}$$

Notice that the presence of congestion costs δ generates an upward-sloping labour supply curve, in contrast to the perfectly elastic labour supply of the original Rosen-Roback framework, where $\delta = 0$. Rearranging (3) also allows to derive an expression of the housing price gradient as a function of relative commuting costs, i.e., of inverse housing demand:

$$q(x) = \left(\frac{\tau(0)}{\tau(x)} \right)^{\frac{1}{\beta}} q(0) \tag{5}$$

As $\tau(0) < \tau(x)$, that is, as the cost of commuting within the CBD is lower than that of commuting from any location x to the CBD, this equation shows that the price of housing at location x decreases with respect to housing prices in the centre ($q(0)$), at a rate that is inversely dependent on the share of housing in expenditure.

Housing production[6] Developers use land $L(x)$, whose rental price is $r(x)$, and building height $h_b(x)$, to produce housing $H(x)$, which is the product of land and height.[7] They have the following

[6]For this portion of the model, I draw inspiration from Glaeser and Gottlieb (2009), also used in Harari (2020).

[7]I make the assumption that developers do not need any workers to produce. This is arguably innocuous even once I take the model to the data, as the construction sector employs about 6% of the workforce nationwide (INSEE, 2023).

profit function, with $\sigma > 1$:

$$\pi_H(L(x), h_b(x)) = Bq(x)L(x)h_b(x) - cL(x)h_b(x)^\sigma - r(x)L(x) \tag{6}$$

The lower σ, the less convex the cost of building housing, hence, the more elastic the intensive margin of housing supply. Then, I deduce the equilibrium on the housing market by maximising (6) and combining it with city-level housing demand, which is the product of worker-level housing demand at location x, $h(x) = \beta w/q(x)$, and population $N(x)$. Denoting $\kappa = \frac{\sigma-1}{\sigma}$ the inverse housing supply elasticity gives the following expression for housing market equilibrium:

$$q(x) = \frac{1}{\tilde{B}} \left(\beta w \frac{N(x)}{L(x)} \right)^\kappa \tag{7}$$

Housing prices at location x decrease with productivity in the construction sector $\tilde{B}$, and increase with the share of housing in expenditure β, wage w and population density $N(x)/L(x)$.

Tradable good production Finally, a representative firm located in the CBD produces the tradable good using all workers N, capital K, and non-traded capital R. Its profit function thus writes as follows:

$$\pi_Y(N, K, R) = (\underline{A}N)^\alpha K^\gamma R^{1-\alpha-\gamma} - r^K K - wN \tag{8}$$

where $\underline{A}$ is labour productivity.[8] I further define it as $\underline{A} = AN^\lambda$, where N^λ captures the labour productivity gains from agglomeration. Later on, A will be negatively affected by local air pollution. Gathering all constant terms in Γ_F and denoting $\tilde{A} = AR^{\frac{1-\alpha-\gamma}{\alpha}}$ to represent local productivity, the first-order and zero-profit conditions yield the following inverse labour demand:

$$w^{1-\gamma} = \Gamma_F \tilde{A}^\alpha N^{-(1-\gamma-\alpha(1+\lambda))} \tag{9}$$

Notice that agglomeration economies λ attenuate the negative effect of hiring an extra worker on nominal wages, because of their positive influence on overall labour productivity.

2.2 Equilibrium relationships at a given spatial surface area

Equations (2), (7) and (9) jointly define an equilibrium vector of endogenous quantities and prices $\{N^*,\ w^*,\ q(0)^*\}$, given the city fundamentals θ, $\tilde{A}$ and $\tilde{B}$ if the city fringe is left to vary. In this

[8] I abstain from modelling the productivity of other factors of production here because I focus on local air pollution, which, if anything, impacts the productivity of capital by a much lower extent than it does the productivity of labour. In related work, Hanlon (2020) makes the assumption that coal use affects productivity as a whole (TFP).

subsection, I start by holding the city fringe, hence the city surface, constant, to solve for equilibrium city size, wages and housing prices. I turn to endogenising it in the next subsection.

Equilibrium city size, S fixed Let us denote the city fringe by $\bar{x}$. Then, city population is equal to the sum of people living in every location x from 0 to $\bar{x}$:

$$N = \int_0^{\bar{x}} N(x)dx = \int_0^{\bar{x}} \frac{H(x)}{h(x)}dx = \int_0^{\bar{x}} \frac{L(x)h_b(x)}{h(x)}dx \tag{10}$$

I use optimal building height $h_b(x) = \tilde{B}q(x)^{\frac{1}{\sigma-1}}$, replace $q(x)$ using housing demand (3) and replace w using labour demand (9). For now, I hold the city's fringe as fixed at $\bar{x}$, and denote $\bar{S} = \int_0^{\bar{x}} L(x)dx = 2\pi \int_0^{\bar{x}} xdx$ the city's (fixed) total surface area. I also denote $T(\bar{x}) = \int_0^{\bar{x}} x\tau(x)^{-\frac{1}{\beta\kappa}}dx$ the accessibility term. Log-linearising and rearranging, equilibrium city population writes:[9]

$$\ln(N^*_{\bar{S}}) = (1-\gamma)\Phi\left[\Gamma_N + \ln(\theta) + \frac{\alpha}{1-\gamma}(1-\beta\kappa)\ln(\tilde{A}) + \beta\ln(\tilde{B}) + \beta\kappa\ln(T(\bar{x}))\right] \tag{11}$$

All three exogenous city fundamentals θ, $\tilde{A}$ and $\tilde{B}$ and temporarily exogenous accessibility term $T(\bar{x})$ enter equilibrium city size positively. Amenities attract workers. Higher productivity also attracts new workers, although feedback through a tighter housing market ($-\beta\kappa$) attenuates this effect. Higher construction productivity lowers housing prices, which attracts workers to the city. But a larger spatial extent has one downside, in that it generates lower accessibility $T(\bar{x})$, which limits the inflow of new residents by creating a disutility.[10]

Equilibrium nominal wages, S fixed Inserting equation (11) into labour demand (9) gives an expression for city-level equilibrium wages.

$$\begin{aligned}\ln(w^*_{\bar{S}}) = \Gamma_w \ &+ \alpha(\delta+\beta\kappa)\Phi\ln(\tilde{A}) \\ &- (1-\gamma-\alpha(1+\lambda))\Phi\left(\ln(\theta) + \beta\ln(B) + \beta\kappa\ln(T(\bar{x}))\right)\end{aligned} \tag{12}$$

Since $\alpha, \beta, \gamma, \kappa, \lambda > 0$, the coefficient associated with labour productivity $\tilde{A}$ is positive: firms in inherently more productive cities are going to pay their workers a higher wage, and even more so in order to compensate them for the congestion costs δ that a larger workforce entails. The signs of the coefficients associated with consumption amenities θ and construction amenities $\ln(\tilde{B})$ are,

[9] In addition to bringing constant terms together into Γs, I denote $\Phi = [(1-\gamma)(1+\delta) - (1-\beta\kappa)\alpha(1+\lambda)]^{-1}$ for easier readability of the equations.

[10] Introducing a simple functional form for commuting costs $\tau(x) = (b+x)^\tau$ as I do in equation (22) and modulo a constant, one may write $\beta\kappa\ln(T(\bar{x}) = \beta\kappa\ln(\bar{S}) - \tau\ln(b+\bar{c})$.

however, *a priori* ambiguous. If the additional labour productivity brought by agglomeration economies $\alpha\lambda$ is large enough to prevail over the income share of the local resource $1-\alpha-\gamma$, then amenities have a *positive* effect on nominal wages. Indeed, amenities attract workers, which depletes wages due to increased labour supply.[11] However, with a larger workforce, agglomeration economies chip in and boost labour productivity, possibly up to the extent that, if λ is large enough, the overall effect of amenities on equilibrium wages is positive. In fact, Section 2.4, in particular equation (16), shows that the empirical existence of an urban nominal wage premium implies that we are in such a situation where $\alpha\lambda > 1-\alpha-\gamma$, such that consumption amenities positively affect equilibrium nominal wages $w^*_{\tilde{S}}$. This implies that correcting for housing costs is necessary to the quantification of cross-city compensating differentials.

Equilibrium housing price, S fixed Then, plugging both equilibrium city size and equilibrium wages into the housing market equilibrium condition (7) renders an expression for equilibrium housing price at the CBD:

$$\begin{aligned}\ln(q(0)^*_{\tilde{S}}) = \ \Gamma_{q_0} \ &+ \alpha\kappa\Phi\left((1+\lambda)\ln(\theta) + (1+\delta)\ln(\tilde{A})\right) \\ &- (1-\alpha\beta\kappa(1+\lambda)\Phi)\left(\ln(B) + \kappa\ln(T(\bar{x}))\right)\end{aligned} \tag{13}$$

Consumption amenities θ raise housing prices, and even more so as agglomeration economies are strong: as seen above, amenities attract workers, and a high enough λ implies that this pushes nominal wages up, thus attracting more workers, and raising housing prices. In a similar way, more productive cities have higher housing prices, even more so if congestion costs δ are high. On the other hand, provided that housing supply is sufficiently elastic (σ low enough), higher construction productivity $\tilde{B}$ has a negative effect on housing prices.

2.3 Incorporating air pollution

Air pollution enters the model in two ways. First, it affects the city's fundamental quality of life θ and its labour productivity $\tilde{A}$. Second, it responds to changes in the city's population N and its surface area S.

Effects of ambient pollution First, I model the fact that local air pollution affects local attributes. Specifically, I consider air pollution Z as a negative shock to both consumption amenities θ (with

[11] The reasoning is similar with respect to the effect of an increase in construction productivity: if it improves, for instance through looser land use regulations, workers face lower housing prices and thus migrate to the city.

elasticity $-z_\theta$) and, perhaps more distinctively, labour productivity (with elasticity $-z_a$):

$$\begin{aligned} \theta &= \underline{\theta} Z^{-z_\theta} \\ \tilde{A} &= \underline{a} Z^{-z_a} \end{aligned} \tag{14}$$

Air quality is widely seen and modelled as a consumption amenity (Chay and Greenstone, 2003; Bayer et al., 2009; Bieri et al., 2023), and I thus assume $z_\theta > 0$. Modelling local air quality as a productive amenity as in the second line of equation (14) is newer in the theoretical part of literature. It is indeed common to model *emissions* of global pollutants as inputs in the production function, which amounts to assuming that they are a by-product of production (Copeland and Taylor, 2004), and, therefore, are in most cases costly to abate.[12] Somewhat in contrast to this, I examine the role played by local *concentration* in air pollution, which is the result of the production process, but is also generated by the transport and residential sectors. In fact, according to French organisation Citepa (2018), 42% of primary $PM_{2.5}$ emissions are from residential heating alone, and 18% originate from road transport. As such, given new evidence on the harmful effects of $PM_{2.5}$ on labour productivity (e.g., Chang et al., 2016), it is reasonable to assume that local air pollution is more of a production disamenity than a production amenity that allows for better firm performance.

$\underline{\theta}$ captures other city-specific amenities, such as the presence of a river or a beach, and $\underline{a}$ captures national labour productivity trends. Note that I make the assumption that local air pollution does not directly affect the third fundamental, construction productivity $\tilde{B}$. In other words, I assume that the potential reaction of land use regulations to local air quality is negligible.[13]

Generating local pollution Second, I model local air pollution Z as a function of city size N and city surface S, with $\phi_N,\ \phi_S > 0$:

$$Z = zN^{\phi_N} S^{-\phi_S} \tag{15}$$

The assumption that city population increase local air pollution also stems from reasons akin to those evoked above: the residential and private road transport sectors account for 52% and 8% of primary $PM_{2.5}$ emissions alone (Citepa, 2018). Adding to this, economic activity is itself concentrated in more populated cities, and it is responsible for 22% of primary emissions, adding

[12] Hanlon (2020) models coal pollution as a by-product of production, i.e., incorporates it in the firm's problem, that also affects productivity. This fits into the historical context of coal being used as an input in British cities during the industrial revolution.

[13] Note that this is also consistent with the fact that the housing sector is assumed not to hire any workers, such that its labour productivity mechanically cannot be affected.

up manufacturing, commercial services, and light commercial transport.[14]

I also assume that a larger spatial extent S allows for a reduction of local air pollution at a given population, for several reasons. First, a larger physical extent allows for the horizontal dispersion of air pollutants. Second, it also leaves more room for green spaces that tend to improve local air quality. I outline a strategy to quantify these two elasticities in Section 4.1, and provide estimates in Section 5.1.

Finally, z represents a local pollution endowment, which captures all city characteristics that may influence its $PM_{2.5}$ levels outside of population and surface. These may range from city topography (e.g., cities located in narrow valleys may experience higher pollution due to less horizontal dispersion) to the pollution intensity of nearby emitters (e.g., cities located close to foreign industrial basins may experience higher pollution due to dominant winds). This can also relate to historical industry specialisation, such as the presence of coal mines or quarries, although the latter are to this day mostly no longer operated in my empirical setting.

2.4 Implications for estimation

Now that the model is exposed for the most part, I draw implications for the empirical estimation of the effects of air pollution on the city-size gradient in nominal wages, housing prices, and real wages. To this end, I introduce subscripts c for cities, and subscript t for time, as I gather panel data for estimation. I focus on the situation where city surface is held constant for now, but I relax this below.[15] Corresponding results focusing on the effect of air pollution independently of its link with city structure are detailed in Appendix A.

Identifying nominal wage and housing price premia The wage premium to agglomeration is measured by regressing (log) wages on (log) employment (density). The first direct implication from the framework developed above is that city size, wages and housing prices are equilibrium objects that are jointly determined given a given level of consumption, production and construction amenities. It immediately follows that there is endogeneity in any regression of the wage-city size type. Following the literature on agglomeration Combes et al. (2010), I instrument city size

[14] In related work, Jégard (2021) shows that due to the failure of firms to internalise air pollution damages, the resulting out-migration of workers implies that large cities are too small. In his model, local emissions are also proportional to city size, but this is obtained only endogenously: as city size grows, there is higher productivity through agglomeration economies, hence higher income, which increases the pollution intensity of firms, and thus air pollution. Hence, city size does affect pollution emissions, but this operates only *via* the production process, since the focus of the paper is to be able to optimally target which firms to tax. Here, I take local air pollution as a result of overall city size, such that it encapsulates all of its vectors and sources, be it production, road traffic or the residential sector.

[15] The qualitative conclusions pertaining to the effects of local air pollution on wage and housing price differentials hold regardless, but should be interpreted with a constant city surface in mind.

using geological characteristics. While I provide further details on the selection of these characteristics in the empirical section of the paper (see Section 4), the model developed here highlights the importance of the nature of the instrumental variable used to generate exogenous variation in city population. Indeed, shocks to labour productivity $\tilde{A}_{ct}$ (resp., construction productivity $\tilde{B}_c$) have a direct effect on wages (resp., housing prices) as an outcome. As such, only a shock on amenities θ_{ct} appropriately identifies the effect of city size on wages or housing prices. In the empirical section, I select geological attributes that reflect soil fertility, as they were historically relevant to location choice, hence to current density *via* location inertia, but do not qualify as current productivity shifters.[16] In the next paragraphs, I provide expressions for the nominal wage premium, housing cost premium and real wage penalty, that a shock on θ_{ct} allows me to identify.

Nominal wage response to city size, S fixed When estimating the urban wage premium by TSLS using an exogenous construction or consumption amenity shifter, one should be estimating:

$$\frac{d\ln(w^*_{\bar{S},ct})}{d\ln(N^*_{\bar{S},ct})} = \frac{d\ln(w^*_{\bar{S},ct})/d\ln(\theta_{ct})}{d\ln(N^*_{\bar{S},ct})/d\ln(\theta_{ct})} = \frac{\eta^W_B}{\eta^N_B} = \frac{\eta^W_\theta}{\eta^N_\theta} = \frac{\alpha(1+\lambda)}{1-\gamma} - 1 \tag{16}$$

Where η^y_x denotes the derivative of y with respect to x. The above equation tells us that an exogenous increase in local population stemming from improved construction amenities affects nominal wages in two ways: higher labour supply does have a direct negative effect on wages, but also generates agglomeration externalities that increase (labour) productivity, and hence, wages (through the $\alpha\lambda$ term). Hence, as long as $\alpha\lambda > 1 - \alpha - \gamma$, that is, as long as agglomeration economies λ are strong enough compared to local costs, we observe a *premium* to density, whereby an amenity-driven rise in population implies an increase in nominal wages. As this is a common observation in practice (Ahlfeldt and Pietrostefani, 2019), I make the assumption that $\alpha\lambda > 1 - \alpha - \gamma$ throughout the rest of this section.

Real wage response to city size, S fixed Using a consumption amenity shock to city size to identify a causal relationship, one should be estimating:[17]

[16] One may argue that soil fertility affects agricultural production, making it somewhat of a labour productivity shifter. This does not present a significant threat to identification in this setting, since in the 21st century, agriculture no longer represents a large part of value added nor employment, particularly in the French context. To alleviate any remaining concerns, I drop the very few agricultural workers (1%) that are recorded in the wage dataset. I also proceed to a careful selection of soil IV in Section **??**.

[17] In this case, as housing prices $\ln(q(0)_{ct})$ are involved as an outcome, the equivalence between using a construction amenity shock B_{ct} and a consumption amenity shock θ_{ct} as in equation (16) breaks, as B_{ct} directly affects $\ln(q(0)_{ct})$ and is therefore not a valid instrument.

$$\frac{d\left(\ln(w^*_{\bar{S},ct}) - \beta \ln(q(0)^*_{\bar{S},ct})\right)}{d\ln(N^*_{\bar{S},ct})} = \frac{d\left(\ln(w^*_{\bar{S},ct} - \beta \ln(q(0)^*_{\bar{S},ct})\right)/d\ln(\theta_{ct})}{d\ln(N^*_{\bar{S},ct})/d\ln(\theta_{ct})} = \frac{\alpha(1+\lambda)(1-\beta\kappa)}{1-\gamma} - 1 \quad (17)$$

This expression closely resembles equation (16), with the addition of the $(-\beta\kappa)$ term. As city increases, nominal wages increase. But this also raises housing demand and housing prices. Whether the urban real wage differential is positive or negative at equilibrium therefore depends on the stickiness of housing supply κ and the share of housing in expenditure β. There is a real wage gain from agglomeration if $\alpha\lambda > \frac{1-\gamma}{1-\beta\kappa} - \alpha$, which is a significantly stronger condition than $\alpha\lambda > 1 - \alpha - \gamma$.[18] Section 5 shows that the reverse does hold on French data, such that larger, denser cities exhibit lower real wages.[19]

Total versus direct nominal wage premium Equation (16) gives the wage premium to density at a given level of air quality, that is, when the increase in city size generates no change in production or consumption amenities. But an exogenous increase in population also raises air pollution (equation 15), which lowers both labour productivity and amenities (equation 14), which in turn affect both city size and wages (equations 11 and 12). An IV regression of (log) wages on (log) city size that does not include air pollution, with city size instrumented by an amenity shifter, therefore identifies the total equilibrium gradient:

$$\frac{d\ln(w^*_{\bar{S},ct})}{d\ln(N^*_{\bar{S},ct})} = \frac{d\ln(w^*_{\bar{S},ct})/d\ln(\theta_{ct})}{d\ln(N^*_{\bar{S},ct})/d\ln(\theta_{ct})} = \frac{\eta^w_\theta - (\eta^w_a z_a + \eta^w_\theta z_\theta)\,\phi_N\,\eta^N_\theta}{\eta^N_\theta\left(1 - \left(\eta^N_a z_a + \eta^N_\theta z_\theta\right)\phi_N\right)} \quad (18)$$

Given that ϕ_N and all η and z coefficients are positive, both the numerator and the denominator of this expression are smaller than their counterparts in equation (16). Whether the pollution channel lowers or raises the measured gradient therefore depends on whether wages or city size react more strongly to amenity and productivity changes at the equilibrium.

The total gradient (18) and the gradient at given air quality (16) are linked by a simple relationship. Given city size and air pollution, wages are pinned down by labour demand (equation

[18] This is true as long as $\beta\kappa < 1$, which is, as abovementioned, a very weak assumption that is easily met in practice. The calibration of the parameters is discussed in Section 6.2. With a share of housing in expenditure β at about 33%, κ would have to be larger than 3.33 for this condition not to hold. Recall that κ is an intensive margin elasticity, such that it being larger than 1 is already a stretch.

[19] This is consistent with the existence of a quality of life premium in larger cities: individuals are willing to give up part of their real wage to live in cities with high quality of life, which makes the latter larger.

9), so that an amenity shock affects wages only through these two variables. It follows that:

$$\frac{d\ln w^*}{d\ln N^*} = \underbrace{\frac{\alpha(1+\lambda)}{1-\gamma} - 1}_{\text{gradient at given } Z} + \underbrace{\left(-\frac{\alpha z_a}{1-\gamma}\right)}_{\partial \ln w/\partial \ln Z}\phi_N \tag{19}$$

where the second term is the direct effect of air pollution on wages at given city size, multiplied by the elasticity of air pollution to city size.

Total versus direct real wage penalty The same reasoning applies to real wages. An IV regression of real wages on city size that does not include air pollution, with city size instrumented by an amenity shifter, identifies:

$$\frac{d\left(\ln(w^*_{\bar{S},ct} - \beta\ln(q(0)^*_{\bar{S},ct})\right)}{d\ln(N^*_{\bar{S},ct})} = \frac{d\ln(w^{R*}_{\bar{S},ct})/d\ln(\theta_{ct})}{d\ln(N^*_{\bar{S},ct})/d\ln(\theta_{ct})} = \frac{\eta^R_\theta - \left(\eta^R_\theta z_\theta + \eta^R_a z_a\right)\phi_N\,\eta^N_\theta}{\eta^N_\theta\left[1 - \left(\eta^N_\theta z_\theta + \eta^N_a z_a\right)\phi_N\right]} \tag{20}$$

where $\phi_N, \eta^R_a, \eta^N_\theta, \eta^N_a, z_\theta, z_a > 0$ but $\eta^R_\theta < 0$. The denominator behaves as in equation (18). In the numerator, the pollution channel has an ambiguous sign. If the productivity effect of air pollution is smaller than its compensating effect, the total real wage gradient is less negative than the gradient at given air quality. If instead the productivity effect dominates, the pollution channel adds to the real wage penalty, and the total gradient is more negative than the gradient at given air quality. As for nominal wages, the two gradients differ by the direct effect of air pollution on real wages times ϕ_N. This requires the instrument to shift amenities θ only, as a construction productivity shifter would affect housing prices directly.

2.5 Allowing for an endogenous city surface area

For now, I have assumed the city's physical surface to be constant. In this context, a positive shock on amenities and labour productivity, such as an improvement in local air quality, triggers an inflow of residents that is fully absorbed through densification. But in this section, I allow for an endogenous reaction of city surface area to changes in local fundamentals, that is, I allow for urban sprawl, and discuss the related implications.

Endogenising the city fringe I use the fact that beyond the city's fringe $\bar{x}$, land is no longer used for housing production, but for agricultural purposes. I start by inserting equation (5) into the housing developer zero-profit condition ($q(x) = \frac{1}{B}\left(r^{(d)}/C\right)^\kappa$), and evaluate this at the CBD, where $x = 0$. I obtain the following land price gradient, which is very similar to the housing price

gradient (5), though with a smaller elasticity to commuting costs:

$$r(x) = \left(\frac{\tau(0)}{\tau(x)}\right)^{\frac{1}{\beta\kappa}} r(0) \tag{21}$$

At this point, it is useful to introduce a functional form for $\tau(x)$. I thus set:

$$\tau(x) = (b + x)^{\tau} \tag{22}$$

with b the very small cost of commuting within the CBD, similarly to Combes et al. (2012).[20] I can thus rewrite the land price gradient as:

$$r(x) = \left(\frac{b}{b + x}\right)^{\frac{\tau}{\beta\kappa}} r(0) \tag{23}$$

Since commuting costs increase with distance to the CBD x, and assuming $\tau > \beta\kappa$, equation (23) shows that land rent decreases as x increases. As such, the city expands up until the point where it is no longer economically sensible for it to, meaning, up until to the point where land rent is equal to the agricultural land rent: at $x = \bar{x}$, I have $r(\bar{x}) = \bar{r}_a$. Combining this with equation (23), using the free-entry condition for housing developers, plugging in equations (3), (9) and (10), and log-linearising, I obtain that the city fringe is equal to:[21]

$$\ln(\bar{x}) = \Gamma_{\bar{x}} + \tilde{\Phi}\left[\alpha\beta\kappa\left((1 + \lambda)\ln(\theta) + (1 + \delta)\ln(\tilde{A})\right) + \beta\left(\alpha\beta\kappa(1 + \lambda) - Q(\sigma - 1)\right)\ln(\tilde{B})\right] \tag{24}$$

This equation shows that cities whose fringe $\bar{x}$ is higher, i.e., that sprawl more, have higher productivity in the construction sector B or tradable sector $\tilde{A}$, and/or higher consumption amenities θ. All elasticities depend negatively on the commuting cost parameter τ (which enters $\tilde{\Phi}$ negatively): the more costly it is for workers to commute from their residence located at x to the CBD, the lower the inflow of workers in reaction to an improvement in amenities or productivity, and thus the more modest the physical expansion of the city.

I then follow the same steps as with a fixed city surface to recover city-level prices and quantities at the spatial equilibrium. Given the complexity of the expressions, I use calibrated parameter values to compute values for each coefficient in Appendix Table A.6. This provides an empirical validation of the discussion of magnitude and sign differences depending on whether S is fixed or free that I give in the next paragraphs.

[20] I refer here to the working paper version of the paper, then superseded by Combes et al. (2019), as the latter published version features a more general model that allows for commuting cost within the CBD to be equal to zero.

[21] Additionally, for simplicity, I denote $Q = (1 - \gamma)(1 + \delta) - \alpha(1 + \lambda)$ and $\tilde{\Phi} = \Phi[\tau - (\tau - 2\beta\kappa)Q\Phi]^{-1}$

Equilibrium city size, S free Going back to the equation for equilibrium city size with a fixed city fringe given in (11), I replace $\bar{x}$ by its equilibrium expression (24), and rearrange to obtain the expression for equilibrium city size with endogenous surface:

$$\ln(N^*) = \underbrace{2(1-\gamma)\beta\kappa\tilde{\Phi}}_{>(1-\gamma)\Phi}\ln(\theta) + \underbrace{\alpha\beta\kappa(2-\tau)\tilde{\Phi}}_{>\alpha(1-\beta\kappa)\Phi}\ln(\tilde{A}) + \underbrace{(1-\gamma)\beta\left(2\beta\kappa-\tau\right)\tilde{\Phi}}_{>(1-\gamma)\beta\Phi}\ln(B) \tag{25}$$

Following a positive shock on amenities, city population increases more when the city fringe is allowed to adjust than when it is not (equation 11), as the upward pressure that is put on housing prices is partially relieved by the expansion of the city onto agricultural land. This is true as long as $\tau < 2\beta\kappa$, that is, as long as the increase in commuting cost does not generate too high a disutility compared to the utility gains brought by new housing supply that pushes prices down. As a consumption and production disamenity that increases with city size but decreases with city surface area, local air pollution enters equation (25) through both θ and $\tilde{A}$, and its effect on equilibrium city size is affected in two opposite ways when the city fringe is allowed to adjust. On the one hand, as the elasticities of city size to θ and $\tilde{A}$ are larger than with a fixed fringe, a given change in air pollution has a larger effect on equilibrium population. On the other hand, a given inflow of workers generates less air pollution, as part of it is accommodated through a larger surface area, which dilutes pollutants. Which of these two effects dominates is an empirical question, which I examine in Section 5.

Equilibrium housing price, S free I follow the same procedure as in Section 2.2 with fixed city surface to obtain an expression for nominal wages, and then housing prices, with an endogenous city fringe.[22]

$$\ln(q(0))^* = \Gamma_{q_0} + Q\tilde{\Phi}\left[\frac{1-2\beta\kappa}{\beta}\ln(\theta) + \frac{\alpha(1-\beta\kappa(2-\tau))}{\beta(1-\gamma)}\ln(\tilde{A}) - (2\beta\kappa-\tau)\ln(B)\right] \tag{26}$$

where $Q = (1-\gamma)(1+\delta) - \alpha(1+\lambda)$, which is positive provided that labour supply is inelastic enough to offset part of the agglomeration gains.[23] The same mechanisms as for nominal wages apply to housing prices: following a positive shock on amenities, the resulting increase in air pollution has a stronger negative effect on housing prices when the city fringe is allowed to react,

[22]The only difference is that I first proceed to adjust the housing supply equation (7) in order to account for the fact that $L(\bar{x})$ is now endogenous. The corrected equation is given in Appendix A.

[23]More specifically, this holds if δ is larger than $\frac{\alpha(1+\lambda)}{1-\gamma} - 1$, which, as equation (16) shows, coincides with the estimand of the urban wage premium. In practice, the urban wage premium is around 0.03, while δ is an order of magnitude higher, and set at 0.4 here.

although the expansion of the city's surface area partly offsets it by dispersing pollutants. Whether this expansion fully offsets the negative effect on housing prices is an empirical question, which I answer below.

3 Data

This section describes all the data sources used throughout the paper. A complete set of summary statistics at the city-level is available in Appendix Table B.3.[24]

3.1 Ambient air pollution

I make use of fine-grained (≈ 1 km $\times$ 1 km) $PM_{2.5}$ concentration data provided by the Atmospheric Composition Analysis Group (ACAG).[25] Satellite measures of aerosol optical depth[26] are used as inputs in a chemical air transport model (GEOS-Chem), so as to retrieve ground-level air pollutant concentration. The estimates are then calibrated to monitor data from the European Environment Agency (EEA), using geographically-weighted regressions (for more details, see van Donkelaar et al., 2021). While Fowlie et al. (2019) suggest being cautious using this data as they may carry measurement error, the authors stress that the main issue is the underestimation of very high concentrations, which does not pose a threat in the comparatively low-pollution context of France. Finally, van Donkelaar et al. (2021) compare their estimates to ground-level monitor measures, and the resulting R^2 is equal to 0.81.

3.2 Wages

Wages are extracted from matched employer-employee panel data (DADS) collected and disseminated by the French Statistical Institute (INSEE).[27] Observations are at the worker-establishment-year level. From 2002 on, the sample is made up of all workers born in October, and thus represents about 1/12 of the population. For each observation, I use information on the individual, namely age, gender, occupation (at the 2-digit level of INSEE's PCS classification), number of hours worked during the year, and wage earnings net of social security contributions, as well as

[24] Like robustness checks, most estimations are also performed using an alternative city definition, commuting zones (CZ), described in this section. Appendix Table D.5 provides the corresponding summary statistics.

[25] The raw data can be found here: `https://sites.wustl.edu/acag/datasets/surface-pm2-5/`.

[26] The presence of dust, sand, smoke and pollutants in the atmosphere absorbs sunlight. Aerosol optical depth (AOD) is a measure of the extinction of sunbeam by these particles within a column of air from the surface of the earth, to the satellite instrument. The higher the value of AOD, the hazier the column of air.

[27] I exploit the panel version of the data so as to be able to introduce worker fixed effects in the first step of my empirical strategy, so as to account for worker sorting across cities.

information on the establishment they work in (municipality, 3-digit sector of the NAF classification). In terms of terminology, (hourly) wage is measured as the ratio of yearly wages to the total yearly number of hours worked during the year. I select 15-65 year-old workers of the private sector, excluding industries with very few observations and agricultural workers, to preserve the validity of the soil instrumental variable. A detailed description of the sample selection is available in Appendix B.1.

3.3 Housing prices

I build a housing price index based on data on the universe of property transactions between 2010 and 2018, available in the DVF database. I select cities for which there are at least 150 recorded transactions every year, such that the main results focus on a sample of 249 cities, observed over 8 years.[28] Available information includes the sale price, transaction date, address, type of property (house, flat, farmland...), and its surface area. I select only house and flat transactions, and thus omit all commercial buildings. In order to account for cross-city differences in the composition of housing stock, I use a housing price index, which corresponds to a city fixed effect introduced in transaction-level regressions. For each of the 8 years of data, I regress the log square meter price of the housing unit on the quarter of the transaction date, a house/flat dummy, and the log of its built surface and the log of the associated plot of land + 1. I follow the model of Section (2) and account for variations in city surface area by controlling for distance to the CBD, such that the housing price index can be interpreted as the housing price at the CBD. I denote this measure $\ln(\widehat{q(0)}_{ct})$ below. I provide more information on sample selection and robustness checks for estimating the housing price index in Appendix B.2.

3.4 Instruments

Soil characteristics I obtain geological information from the European Soil Database (ESDB), maintained by the European Soil Data Centre (ESDAC). It provides information on biological, chemical, hydrological and mechanical soil properties, disseminated as rasters of 1 km by 1 km. I select instruments that are a) not derived from current human activity and b) both theoretically and quantitatively correlated with density and surface area, as discussed in Section 4.3 and Appendix C.1. The first selected instrument is the cation exchange capacity (CEC) of the subsoil, the second is the depth to a gleyed soil horizon (or layer), and the third is the organic carbon content

[28] Due to historical reasons, the three departments of Alsace-Moselle are excluded from the scope of the data, such that cities in this region must be excluded from the estimating sample. Wage results are robust to expanding the sample to the full set of 310 cities over 17 years (2002-2018), as shown in Appendix.

of the topsoil.

In the raw raster data, pixels take categorical values, with, e.g., 0 indicating a low level of subsoil CEC, 1 indicating a medium level of subsoil CEC. I compute the share of a city's surface area where subsoil CEC is medium, the share where the depth to a gleyed horizon is medium, and the share where topsoil organic carbon content is high or very high. Panels (a) to (c) of Appendix Figure C.1 map the three instruments, which can be compared with a map of employment density in Panel (d).

Planetary boundary layer height The planetary boundary layer height (PBLH) instrument is constructed using the ERA5 dataset of the Copernicus European project, a gridded dataset of 0.25° $\times$ 0.25° (approximately 26 km $\times$ 26 km).[29] The value of the planetary boundary layer height in metres is provided at an hourly frequency. I first aggregate the gridded data up to the daily level, for all days from 01/01/2002 to 31/12/2018. Second, I proceed to compute the mean daily PBLH at the city level. The instrument captures the fact that the concentration in fine particles increases when the PBLH is low, and is therefore constructed as follows. Denoting $\overline{\text{PBLH}}_c$ the mean city-specific PBL height over all days of the study period, the dummy LowPBLH_{cd} takes the following value for city c on day d:

$$\text{LowPBLH}_{cd} = \mathbb{1}\left\{\text{PBLH}_{cd} < a\,\overline{\text{PBLH}}_c\right\}$$

Then, I sum this dummy over all days during year t to obtain the number of days where the PBLH was *unusually* low during the year. The more days with a low PBLH, the higher the yearly mean concentration in $PM_{2.5}$. The main results are given using a factor of $a = 0.7$ to consider a day as a low-PBLH day, but the results are robust to using factors of 0.5 to 0.9.[30]

Historical population I use past population and past population in the central municipality of the city to instrument for city total surface area. Municipality-level historical population data is drawn from the INSEE-EHESS-Cassini project.[31] The first census started in 1791, but a large proportion of municipalities are missing. I use 1806, the first most complete census year, and 1836.

[29] I use the ERA5 dataset (Hersbach et al., 2018), downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store (CDS). DOI: `10.24381/cds.adbb2d47`

[30] Taking more extreme values of a does not significantly change the point estimates, but induces a strong loss in precision, as the number of selected days is mechanically either very low (if below 0.6) or very high (if above 0.9), with low variation across cities, and thus, less statistical power.

[31] Each municipality has its own webpage, where one can find its population for each census year from 1793 to 1999: `http://cassini.ehess.fr/fr/html/`.

3.5 Control variables and alternative outcomes

Weather controls As for PBL height, meteorological variables are extracted from the ERA5 dataset of the Copernicus European project. I use hourly measures of surface temperature, wind speed and total precipitation, which I average over the whole city surface area. I then aggregate wind speed to the yearly level, and temperature to the season level (summer and winter) by taking the simple average, for each year from 2002 to 2018. I aggregate the precipitation data by taking the sum over the year.

Geographical controls I also control for the presence of different amenities, namely mean elevation, ruggedness, the share of the surface covered by mountains, and dummies capturing access to economic and environmental amenities (waterway and the sea), using data from BD-TOPO, which is produced by IGN, the French National Institute of Geographic and Forest Information, and Observatoire des Territoires.

Measure of labour productivity I approximate labour productivity through (raw) mean value added per worker. To construct this proxy, I use information extracted from firm-level balance sheet data (in the FARE and FICUS databases). Due to the fact that the information is given at the firm level, as opposed to the establishment level, and that I need to be able to pinpoint exactly where each establishment is located in order to avoid wrong inference, I focus on private-sector single-establishment firms. I compute value added per worker as the ratio of gross value added to total employment in the firm, which I then aggregate to the city level.

3.6 City definition

Cities are defined here based on the French Statistical Institute (INSEE) definition of metropolitan areas (MSA, *aires urbaines* in French). This partitioning was built in 2010 using 2006 Census data[32], and defines a city as a cluster of municipalities that constitute an urban pole and its suburban ring, accounting for both inner and outer suburbs. More formally, municipalities are grouped into the same MA if they a) exceed a certain population threshold, b) constitute a built-up continuum, and c) comprise a minimum of jobs, which generate sufficient within-MSA commuting flows. It thus takes into account the location of buildings and economic activity, and commuting patterns, which are all important determinants of both air pollution and wage levels. In terms of sample selection,

[32]Until 1999, INSEE ran exhaustive censuses. In 2004, INSEE switched to annual surveys, covering all municipalities with less than 10,000 inhabitants and 40% of dwellings in municipalities with more than 10,000 inhabitants over a five-year period. Hence, the "2006 Census" refers to the wave of data collected over the 2004-2008 period.

I focus on urban areas that have at least 10,000 jobs, and drop rural areas and small towns. My final sample of cities represents 45 million people (75% of the country's population).

An additional set of results, provided in Appendix **??**, relies on commuting zones (CZ, or *zones d'emploi* in French), also drawn by INSEE in 2010. This zoning divides country based solely on commuting flows, which makes it a relevant definition of local labour markets, in such a way that it covers the entire territory. It also allows for easier comparison with the earlier literature on agglomeration economies in France, but it is less directly relevant to study air pollution, and often features several CBDs, thus preventing a clean mapping from the model to its empirical implications.

4 Identification strategy

This section turns to the translation of the above theoretical set-up into reduced-form, estimable equations. First, it focuses on identifying the effect of city structure on local air pollution, and then addresses the impact on nominal and real wages.

4.1 Density and air quality: Identifying ϕ_N and ϕ_S

The model draws attention to the fact that the wage gains from agglomeration may be mismeasured, if city size and/or surface area does induce higher air pollution, that is, if $\phi_N,\ \phi_S > 0$ in equation (15). This is indeed a necessary condition for the urban wage premium equation (16) not to be equivalent to equation (18), and for the real wage differential equation (17) not to be equivalent to equation (20). The first step of the empirical analysis therefore consists in verifying that the previously identified relationship is not specific to the United States or Germany, but exists in the French context. To this end, I estimate the empirical counterpart of equation (15), which writes:

$$\ln(\mathrm{Z}_{ct}) = \phi_N \underbrace{\widehat{\ln(\mathrm{density}_{ct})}}_{\ln(N_{ct})/\ln(S_c)} - \tilde{\phi}_S \underbrace{\widehat{\ln(\mathrm{area}_{ct})}}_{\ln(S_c)} + \phi_\Gamma \Gamma_{ct} + \eta_t + \varepsilon_{ct} \tag{27}$$

The outcome is log $\mathrm{PM}_{2.5}$ concentration, weighted by neighbourhood population. $\ln(\mathrm{density}_{ct})$ is city density, defined as city-level density of employment (or population, in robustness checks). As I control for (instrumented) land area, the ϕ_N coefficient can be interpreted as both the effect of a 1% increase in density, or a 1% increase in city size.[33] Γ_{ct} includes both time-varying weather

[33] In such a case, given the log specification, I deduce that $\tilde{\phi}_S = \phi_N - \phi_S$, such that coefficient $\tilde{\phi}_S$ on log area is not directly interpretable as the effect of city surface, but should be computed as the difference between $\hat{\phi}_N$ and $\hat{\tilde{\phi}}_S$. In alternative specifications, I also use $\ln(\mathrm{size}_{ct})$, defined as $\ln(N_{ct})$ directly.

controls, and time-invariant geographical controls, as they may affect both local air pollution levels and population density. More specifically, I control for average wind speed, total precipitation and mean temperature, differentiating between summer and winter. Geographical controls include dummies for access to the sea and the presence of a waterway, as well as the share of land covered by mountains, mean elevation and ruggedness.

Even conditionally on geographical controls, employment in a given city is likely endogenous, as households have a preference for cleaner air, and would move away from a very polluted city to a cleaner one. The model shows that, as a consumption disamenity, air pollution negatively affects workers' willingness to locate in polluted cities, thus negatively affecting the latter's density. As a production disamenity, it reduces firms' incentive to locate in polluted cities, also reducing employment. This induces a reverse causality bias in the OLS estimation, which would result in an underestimation of ϕ_N. As evoked in the theoretical section, I instrument employment density using three soil characteristics that determined agricultural potential: subsoil cation exchange capacity and topsoil organic carbon content, which relate to soil fertility, and the depth to a gleyed horizon, which relates to drainage capacity, in the spirit of Rosenthal and Strange (2008b). Indeed, when agriculture was a key part of the economy, these characteristics were a strong determinant of household location decisions, and therefore of today's spatial distribution of workers, but they are unlikely to affect local air pollution other than through city structure (see Appendix C.1).[34] Moreover, I use log 1806 population and 1836 central-municipality density to instrument for city surface area. Historically populated cities have expanded more into space on average (Combes et al., 2024), but this is exogenous to current pollution conditionally on controlling for the presence of long-lasting amenities such as access to the sea, or current-day weather amenities. The first-stage regressions, for city size and surface area respectively, thus write as follows:

$$\begin{aligned} \ln(\text{density}_{ct}) &= \pi_0^D + \pi_{IV}^D(\text{soil})_c + \pi_\Gamma^D \Gamma_{ct} + \pi_t^D + \epsilon_{ct}^D \\ \ln(\text{area}_{ct}) &= \pi_0^S + \pi_{IV}^S \ln(\text{lagged pop.})_c + \pi_\Gamma^S \Gamma_{ct} + \pi_t^S + \epsilon_{ct}^S \end{aligned} \tag{28}$$

where Γ_{ct} include geography and weather controls and the excluded instruments of the other endogenous variables.

[34] Sand, which is correlated with soil fertility, is technically included in PM_{10}. But these are particulates that are coarser than $PM_{2.5}$, so this does not pose a threat to the causal identification of the coefficients. I provide a more thorough discussion of the nature and relevance of these instruments for city structure in Section 4.3 on the identification strategy for the main wage and housing price equations.

4.2 Main wage and housing price equations

The second and main step of the reduced-form analysis brings air pollution into the city size-wage and city size-housing price relationships, and estimates the following equation:

$$\ln(Y_{ct}) = \beta_0 + \beta_{\text{N/S}} \ln(\text{density}_{ct}) + \beta_{\text{Z}} \ln(\text{Z}_{ct}) + \beta_S \ln(\text{area}_c) + \beta_{\Gamma} \Gamma_{ct} + v_t + u_{ct} \tag{29}$$

where $\ln(Y_{ct})$ is either the log of mean wage, or the log of the housing price index $\ln(\widehat{q(0)}_{ct})$. Log mean wage has three definitions:

- W_{ct}^{raw}: the raw mean of log wages at the city×year-level
- $\hat{W}_{ct}^{adj}$: the mean of log wages, correcting for local sectoral composition and worker fixed effects, following equation (39) below
- $\hat{W}_{ct}^{adj} - \beta \ln(\widehat{q(0)}_{ct})$: real wages, defined as nominal wages corrected for housing costs (housing price index multiplied by the share of housing in expenditure)

Γ_{ct} includes the same time-varying weather controls (wind speed, precipitation, winter and summer temperature) and time-invariant geographical controls (sea, waterway, mountains, mean elevation, ruggedness) as above. The estimated equations also control for time trends, the empirical counterpart of constant utility v_t. My main coefficient of interest is $\beta_{N/S}$. Its interpretation depends on whether air pollution is included in equation (29). When it is not, $\beta_{N/S}$ identifies the total gradient of wages with respect to density, which combines agglomeration economies λ and the effect of the air pollution that density generates (equation 18 for nominal wages, equation 33 for housing prices, equation 20 for real wages). When instrumented air pollution is included, $\beta_{N/S}$ identifies the gradient at a given level of air quality, and β_Z the direct effect of air pollution. The two are linked by $\beta_{N/S}^{\text{total}} = \beta_{N/S} + \beta_Z \, \phi_N$, where ϕ_N is the elasticity of $PM_{2.5}$ to density estimated in equation (27) with the same instruments for density and surface area. This relationship should not be read as an omitted-variable correction in the OLS sense: both coefficients are IV estimands. It holds because, given density, surface area and air quality, wages are determined by labour demand, so that the instruments for city structure affect wages only through these variables. It requires that these instruments affect $PM_{2.5}$ only through density and surface, which I discuss in Section 4.3 below.

The model yields two predictions. For nominal wages (W_{ct}^{raw} or $\hat{W}_{ct}^{adj}$), both disamenity channels lower wages, so that $\beta_Z \phi_N < 0$ and the gradient at given air quality should exceed the total gradient. For real wages ($\hat{W}_{ct}^{adj} - \beta \ln(q(0)_{ct}^V)$), the sign of $\beta_Z \phi_N$ is ambiguous: it is negative if the productivity effect of air pollution dominates the compensating effect, positive if the compensating effect dominates, and zero if the two cancel out.

4.3 Tackling endogeneity issues

Equation (29) is subject to three main endogeneity issues that invalidate the identification of coefficients $\beta_{\mathrm{N/S}}$, β_{S} and β_{Z} by OLS. Two main mechanisms threaten the causal identification of the first two city structure coefficients. First, as highlighted by the model, more productive places attract more workers, making them larger and denser, while also affecting wages. Second, some workers may have a stronger preference for density than some other workers due to different preferences for amenities induced by density, thus introducing unobserved heterogeneity in the wage level observed in the city. The coefficient associated with air quality, β_Z is also endogenous due to it being both a consumption disamenity and a production disamenity: lower perceived amenity levels and lower productivity from Z have an effect on both wages and employment. I discuss how I tackle these issues below.

Instrumenting for city structure: soil characteristics As the model shows, a valid instrument for city size N_{ct} or surface S_c must shift consumption or construction amenities (θ_{ct} or $\tilde{B}_c$), not firm productivity $\tilde{A}_{ct}$, since a productivity shock affects wages directly and would violate the exclusion restriction. When housing prices (and hence real wages) are the outcome, the instrument must shift consumption amenities θ_{ct} alone, as a construction-amenity shock $\tilde{B}_c$ affects housing prices directly. I therefore instrument employment density with city-level soil characteristics, which should affect wages only through density. More specifically, I rely on a persistence argument: soil characteristics determined agricultural potential, and hence settlement, when agriculture dominated the economy, and settlement patterns persist through the inherited housing stock and local networks. I also select instruments that are not derived from current human activity, and thus exclude direct measures of land use and erosion.

The first set of instrumental variables pertains to soil fertility. The first instrument is subsoil cation exchange capacity (CEC), a chemical property of the lower layers of the soil. As most of the main plant nutrients are cations (e.g., calcium Ca^{2+} or magnesium Mg^{2+}), the higher the quantity of cations that can be retained in the soil, the higher its fertility. The second instrument is medium depth to a gleyed horizon of the soil. In a gleyed layer, water replaces oxygen, such that, when the gleyed layer is too shallow, plants lack oxygen to grow. On the other extreme, when the layer is too deep, soils tend to lack water, although it is also necessary for plant growth. Hence, medium depth to gleyed soils tends to be associated with better drainage capacity, favouring plant growth. In the same spirit, the third instrument is topsoil organic carbon content (OCC), which also helps water retention and nutrient circulation. The maps of Appendix Figure C.1 allow to see the degree of correlation between medium subsoil CEC (Panel a), medium depth to a gleyed horizon (Panel

b), high topsoil organic carbon content (Panel c) and current-day employment density (Panel d). They also show that the three IVs are not very correlated.[35]

I introduce these instruments using quadratic terms, as I posit that city size and surface area may not be best predicted by a linear combination of these three instruments, but rather, that the relationship may be inverted U-shaped. Soils had to be fertile enough to support early settlement, but not so fertile as to lock local economies into agriculture, with lower current-day population. Appendix Figure C.2 plots city size and surface against each of the three soil instruments and does show a hump-shaped relationship. Introducing a quadratic term thus theoretically increases the relevance of the soil instruments. First-stage regressions in Table C.1 for city size and Table C.2 for surface area confirm that these features are indeed strong determinants of city structure, even when controlling for geographical factors such as hilliness, and that the relationship is indeed inverted U-shaped.

The exogeneity of these instruments to present-day wage and housing price outcomes seems satisfied. Given the current weight of agriculture in the French labour market and GDP (about 2% in both cases), the likelihood of soil fertility affecting current productivity and wages is unlikely. Still, I exclude agricultural workers from the analysis to preserve exogeneity. Appendix C.1 further discusses the validity of each separate soil instrument.

Instrumenting for city structure: Lagged population The second set of instruments relies on the persistence argument more explicitly, and thus relies on long lags of population. I take two measures: 1806 total population within the current-day boundaries of the city, and 1836 density in the central municipality of the city. The first instrument is a "pure" historical lag. It is relevant because cities with a larger historical population have experienced larger absolute increases in population over the following two centuries, and these increases were mostly accommodated by an expansion of the urbanised area rather than by a rise in density in the historical core (Combes et al., 2024). A larger population in 1806 is therefore associated with a larger surface area of the metropolitan area today. The second instrument rests on a complementary idea: cities whose central municipality was already densely populated in 1836 had less scope for further densification of their core, such that subsequent population growth was more likely to take place through outward expansion. These settlement patterns are, however, likely exogenous to present-day productivity, wages, or housing prices. In particular, I choose pre-industrial census years so as to avoid any concerns related to past city size being correlated with current-day industrial specialisation.

[35] All coefficients of correlation between the three IVs are low. The coefficient of correlation between subsoil CEC with medium depth to gleysoil is 0.184, and with topsoil OCC 0.074. The coefficient of correlation between the medium depth to gleysoil instrument and the topsoil OCC instrument is 0.093.

Correcting the wage measure I use matched employer-employee data to construct wage measures. A first set of results uses raw log individual wages aggregated at the city×year-level W^{raw} as the outcome. Because skilled workers sort into denser cities (Glaeser and Mare, 2001; Diamond, 2016), especially in France (Combes et al., 2010), this raw measure risks overestimating the density coefficient. My main set of results therefore uses an *adjusted* wage W^{adj}, purged of unobserved worker sorting and localisation economies via the two-step method of Combes et al. (2008): exploiting the individual panel structure, worker fixed effects absorb unobserved worker heterogeneity, and I recover the adjusted wage as a city-year fixed effect that nets out individual and sector-city composition, as written in Appendix equation 39. Computing this adjusted measure before estimation, rather than adding controls to equation (29), is particularly important here because it accounts for correlations between city-sector composition and unobserved city characteristics: some cities plausibly hold a partly unobserved comparative advantage in polluting economic activity, inducing a sectoral specialisation correlated with air pollution that a specification based on controls would not properly absorb. I discuss sample restrictions and the estimation of city-level adjusted wages in further detail in Appendix Section B.1.

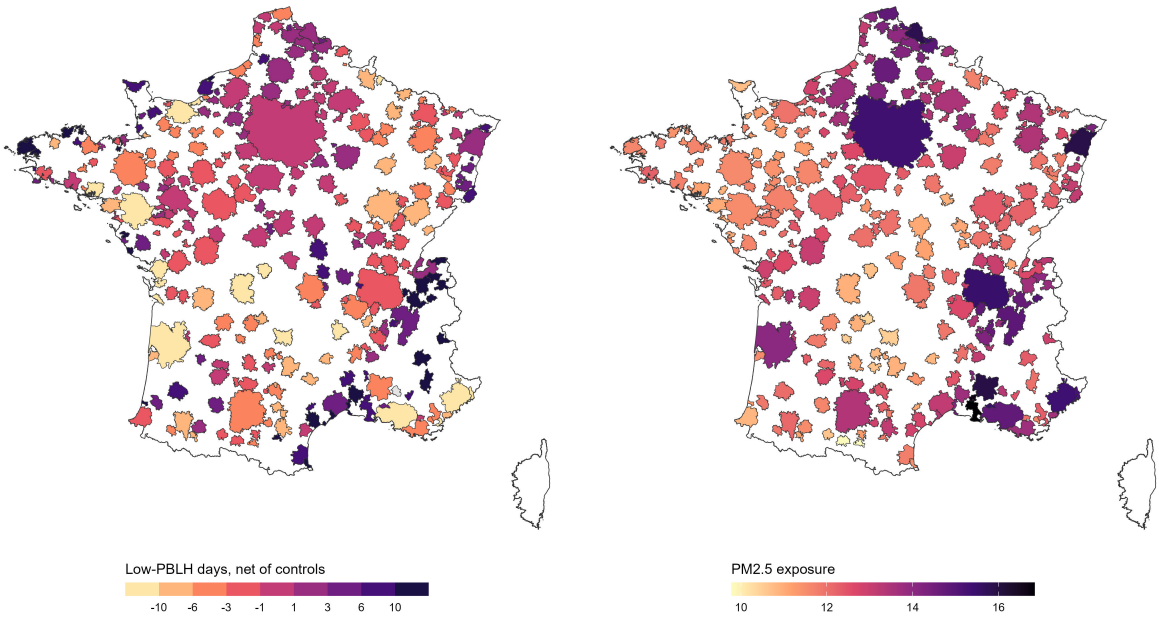
Instrumenting for air pollution The inclusion of $PM_{2.5}$ in the wage and housing price premium regressions introduces another source of endogeneity. Even controlling for (instrumented) city structure, remaining factors, in particular unobserved amenities such as green spaces, may be correlated with both air quality and wages and housing prices.

I address this by exploiting exogenous variation in the height of the planetary boundary layer (PBL), the lowest layer of the atmosphere and the one in which we live. Its height varies between roughly 100 m and 3,000 m depending on weather conditions, and the lower it sits, the less room pollutants have to disperse vertically, raising ground-level particulate concentrations.[36] Because this height depends on topography, some cities sit chronically lower than others, so the yearly mean PBL height could itself be endogenous to the presence of other amenities like mountains. I therefore compare each city against its own long-term mean rather than a common one: for each city×year, the instrument counts the days on which the PBL falls more than $1 - a = 30\%$ below the city's own long-term expectation ($a = 0.7$). This comparison captures how often a city departs from its own norm rather than its fixed topographic position. Appendix C.2 shows that a common benchmark would instead recover that fixed position, gives justification behind the choice of a,

[36] In the data, yearly city averages range from 282 m to 676 m, with a mean of 550 m, and daily values range from 9.6 m to 2,213 m.

Figure 1: Low PBL height instrument and $PM_{2.5}$ exposure

(a) Days with a low PBL height, residualised (b) $PM_{2.5}$ exposure



Sources: ERA5, ACAG, INSEE, IGN.
Notes: All values are averages over the 2006-2018 sample period. Panel (a) residualises the low PBL height instrument using the set of geography and weather controls used in main regressions. Panel (b) shows $PM_{2.5}$ exposure, measured by taking a population-weighted average of $PM_{2.5}$ concentration at the neighbourhood level. This is the main measure of air pollution used throughout the paper.

and provides evidence on the absence of correlation between the PBL height instrument and the soil fertility instruments. I control for the long-term mean PBL height in the main regressions, so that systematic long-term differences in PBL height are not confounded with the short-term cross-city variation I exploit.

Figure 1 maps the mean value of the (residualised) PBL height instrument alongside mean $PM_{2.5}$ concentration. Residualising shows that, after controlling for weather and geography, substantial variation in the number of low-PBL days remains across cities, which is what I exploit; the two panels also line up, with high values of both in the North, the Rhône valley and the Alpine foothills, lower values in the East, and variations elsewhere.[37]

A remaining concern is that a low PBL raises the concentration of other pollutants as well, which would violate the exclusion restriction if the PBL affected them other than through their

[37] Appendix Figure C.4 shows the raw PBL height instrument in Panel (a), the raw $PM_{2.5}$ measure in Panel (b), and the main $PM_{2.5}$ measure, which weights air quality by neighbourhood pollution and is also shown in Panel (b) of Figure 1, in Panel (c).

Table 1: Raw correlations between instruments for city structure and air pollution

Instrument	Soil features			Historical pop.	
	% mid gleysoil depth	% high CEC	% high OCC	log 1806 pop.	log 1836 central density
PBL height	-0.139	-0.026	-0.102	-0.035	-0.064

correlation with $PM_{2.5}$. At the between-city level of this analysis (as opposed to the within-city or neighbourhood level) cross-city differences in $PM_{2.5}$ are likely a sufficient proxy for differences in overall air quality.

4.4 Assessing instrument validity

Strength of the instruments Table C.1 allows to assess the strength of the instruments for city size. Columns (1)-(3) provide results for the first-stage regressions when pollution is the outcome (equation 28), or when mean wage is used as an outcome but pollution is not yet included. Columns 4-6 apply when wages are used as an outcome (real wages in Column 6) and $PM_{2.5}$ is also instrumented (equation 12). Soil features are strong determinants of employment density, and remain so conditional on geographical controls and the PBL instrument for local air pollution.

Table C.3 gives first-stage regression results for $PM_{2.5}$ concentration. The number of days with an unusually low planetary boundary layer height is indeed a strong determinant of yearly pollution levels, even when controlling for weather conditions (F-statistic of 39.7 in the main regression in Column 6).

Independence of the instruments In equation (29), density, surface area and $PM_{2.5}$ concentration are instrumented simultaneously. Identification requires that the instruments for city structure and for air pollution provide independent sources of variation, that is, that they do not move the endogenous variables in proportional ways. Table 1 shows that the raw correlations between the PBL height instrument and the instruments for city structure are all below 0.15 in absolute value. I further regress the PBL height instrument on each instrument for city structure, conditional on the controls of the main specification. Appendix Figure C.5 shows the resulting coefficients, in bold purple when the low-PBL day threshold is set at 30%, as in the main regressions, and in orange for other thresholds. None is statistically significant, which indicates that the instrument for air pollution is unrelated to the instruments for city structure, conditional on controls.

5 Reduced-form results

This section focuses on reduced-form estimations of the main relationships. First, it provides and discusses new estimates of the effect of city density and city surface area on local air pollution levels. Second, it provides evidence on the joint effect of local air pollution, city population and city surface area on nominal wages, housing prices, and, in turn, real wages. It highlights how the urban wage and housing cost premia respond to controlling for local air pollution.

5.1 City structure and air quality: Estimating ϕ_N and ϕ_S

Table 2 shows the estimation results of the effect of population density and surface area on local $PM_{2.5}$ concentration, both for OLS and TSLS. The coefficients shown here are therefore provides estimates for ϕ_N and ϕ_S in the pollution generation equation (15), with ϕ_N being equal to $\beta_{N/S}$ when area is controlled for (Columns 3, 5-6), and $\phi_S = \beta_S - \beta_{N/S}$. Results consistently show that higher population density increases $PM_{2.5}$ concentration. The downward bias of the OLS estimates (Columns 1-3) is consistent with selective sorting across cities, whereby workers would flee pollution that negatively affects amenities θ_{ct} and productivity $\tilde{A}_{ct}$. Formally, a 1% increase in density results in a .12% increase in local air pollution. For the average city in the sample, i.e. for a density of 196 workers by km^2, doubling density raises $PM_{2.5}$ concentration from 12.18 to 13.64 $\mu g/m^3$. This roughly corresponds to a jump from the median to the third quartile of the cross-city distribution of $PM_{2.5}$ concentration. Turning to the effect of city surface area, we need to deduct the coefficient associated with density from that associated with surface area to deduce the actual value. This is shown in the row right below the main estimates. I find that holding city population fixed, city surface has a negative effect on $PM_{2.5}$ concentration. This is attributable to the fact that a larger spatial extent allows for pollutants to disperse horizontally. It is also associated with more green space (Kahn and Walsh, 2015), which was itself shown to decrease $PM_{2.5}$ concentration (Ai et al., 2023).

An interesting finding here is that $\hat{\beta}_S$ is roughly equal to zero, such that $\Phi_N \approx -\Phi_S$. In other words, a one percent increase in population raises local air pollution by about as much as a one percent expansion of the city's spatial extent lowers it. Sprawl neutralises the air pollution effect of a higher population to the extent that surface area grows as fast as population. The pollution channel thus operates on the density margin rather than on the sprawl margin, so that the negative effect of pollution on wages and the capitalisation of pollution into housing prices should be read as costs of density rather than total size *per se*.

Table 2: Impact of city structure on air pollution

	Outcome: (log) $PM_{2.5}$					
	(1)	(2)	(3)	(4)	(5)	(6)
(log) density	0.060***	0.051***	0.051***	0.152***	0.160***	0.115***
	(0.007)	(0.006)	(0.007)	(0.037)	(0.040)	(0.027)
(log) area			0.0006		0.012	0.006
			(0.004)		(0.008)	(0.007)
Implied $\hat{\phi}_S = \hat{\beta}_S - \hat{\beta}_{N/S}$			-0.050		-0.148	-0.109
Density instrumented?				X	X	X
Area instrumented?					X	X
Year fixed effects	X	X	X	X	X	X
Geo and weather controls		X	X	X	X	X
Observations	4,628	4,628	4,628	4,628	4,576	2,214
Cities (clusters)	356	356	356	356	352	246
Mean outcome	2.51	2.51	2.51	2.51	2.51	2.48
F-stat, (log) density				36.26	30.90	29.39
F-stat, (log) area					583.55	347.04
Kleibergen-Paap				36.42	28.32	29.59

Notes: Standard errors clustered at the city (MSA) level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$
$PM_{2.5}$ measured as (municipal-)population-weighted $PM_{2.5}$ at the city scale. Geography controls include mean elevation, ruggedness, share of land covered by mountains, and dummies for sea access and waterway presence. Weather controls include summer temperature, winter temperature, mean wind speed and total precipitation. IV for density and surface area: 1806 population, 1836 density in the central municipality of the current city, soil characteristics.

Discussion on magnitudes These estimates of the elasticity of $PM_{2.5}$ to density are very close to those obtained by Carozzi and Roth (2023), who find values ranging from 13% to 15%, in the U.S. context. But, they are larger than the estimates of Borck and Schrauth (2021) on German data for PM_{10}, a category that includes $PM_{2.5}$ and coarser particulates, as those lie around 6-8%. One may *a priori* hypothesise that differences between estimates drawn from German data and estimates drawn from U.S. data arise due to European cities being characterised by a lower reliance on individual cars, and wider public transport networks (Blaudin de Thé et al., 2021; Conwell et al., 2023). If, for a given increase in density, European cities have a larger response in terms of public transit development and take-up, then this can translate into smaller increases in local pollutant concentrations than in U.S. cities. Table 2 comes to contradict this theory, as estimates for France are closer to the U.S. than Germany.

I find that instead of differences in context, differences in city definition are what appears to drive the gap in these estimates. Carozzi and Roth (2023) use a definition that is close to this paper's, in that they use Metropolitan Statistical Areas (MSAs), as delineated by the US Census

Bureau.[38] On the other hand, cities in Borck and Schrauth (2021) are defined as districts and labour market regions. This comes closer to the definition of "commuting zones" (CZ) introduced in Section 3. Performing the same analysis on this alternative observation level gives the results shown in Appendix Table D.6.[39] These results are comparable, both in terms of city definition (CZ) and magnitude, to the 6-8% range of Borck and Schrauth (2021). These discrepancies can be explained by the fact that the MSA zoning drops rural areas, which are less densely populated. As such, taken together with the convexity of the effect of density on $PM_{2.5}$ concentration, one obtains larger estimates for MSAs than for CZs. This exercise thus suggests that instead of inherent differences between countries, differences in estimated ϕ_N rather pertain to non-linearities in the relationship between density and air quality.

Mechanisms Table D.8 provides suggestive evidence of the mechanisms through which higher density hinders air quality. Higher density is associated with both higher total residential emissions and higher total traffic emissions, which is expected. While it is associated with lower traffic emissions per capita (per worker in Column 3) and per car commuter in Column 4), the overall distance travelled by car is larger (Column 5). More importantly, the average distance travelled by car commuters decreases with density, which means that the reaction in terms of public transport development or usage is not sufficient, even for short distances. This is also due to city surface: at a constant population, cities with a larger spatial extent have lower total emissions, but the distance that workers have to commute is much larger (Columns 5 and 6). Overall, there is a reaction in per capita terms, but it is not sufficient to offset the total size effect, and emissions are higher.

5.2 Air quality and the urban wage premium

As established right above, city structure, in particular density, affects local air pollution, such that part of the wage and housing-cost differences observed across cities of different sizes operates through air quality rather than through city size alone. This subsection estimates how accounting for this channel affects wage, housing price, and real wage differentials to agglomeration. It discusses the pollution-only estimates, reported in Appendix, as a consistency check.

Attenuation of urban wage premium Table 3 shows the main results on the urban wage pre-

[38]The US Census Bureau defines an MSA as a "core area containing a substantial population nucleus, together with adjacent communities having a high degree of economic and social integration with that core". The French Statistical Institute provides a similar definition in its own MSA zoning, as mentioned in Section 3.6.

[39]Figure D.1 in Appendix gives maps of the values of the soil characteristics used as IV (Panel a) and 2010 density (Panel b) for French commuting zones. As one can see on these maps, urban areas do not cover the entire territory, unlike employment zones, but still expand to a large part of the country, as they include outer suburban municipalities.

Table 3: Local air pollution and the urban wage premium

	Outcome: (log) nominal wage W^{adj}					
	(1)	(2)	(3)	(4)	(5)	(6)
(log) density	0.026***	0.040***	0.014	0.031***	0.070**	0.044***
	(0.005)	(0.003)	(0.012)	(0.006)	(0.035)	(0.011)
(log) area		0.029***		0.030***		0.032***
		(0.002)		(0.002)		(0.002)
(log) $PM_{2.5}$					-0.675***	-0.129*
					(0.239)	(0.074)
Implied $\hat{\eta}_S = \hat{\beta}_S - \hat{\beta}_{N/S}$		-0.011		-0.001		-0.012
Density instrumented?			X	X	X	X
Area instrumented?				X		X
$PM_{2.5}$ instrumented?					X	X
Year fixed effects	X	X	X	X	X	X
Observations	2,241	2,241	2,241	2,232	2,223	2,223
Clusters (cities)	249	249	249	248	247	247
Mean outcome	0.523	0.523	0.523	0.523	0.523	0.523
F-stat, (log) density			30.57	41.68	38.26	33.60
F-stat, (log) area				367.16		299.69
F-stat, (log) $PM_{2.5}$					34.00	37.43
Sargan, p-value			0.001	0.003	0.001	0.154

Notes: Standard errors clustered at the city (MSA) level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$
Geography controls include mean elevation, ruggedness, share of land covered by mountains, and dummies for sea access and waterway presence. Weather controls include summer temperature, winter temperature, mean wind speed and total precipitation. IV for density: soil characteristics. IV for area: 1806 population. IV for $PM_{2.5}$: number of days with an unusually low PBL height.

mium, that is, the nominal wage premium to density. Note that the outcome variable is W^{adj}, such that the potential confounding effect of localisation economies and unobserved worker heterogeneity is removed. First, estimates in Columns 1-4 confirm existing estimates of nominal wage gains from agglomeration, including recent ones for France (Berson et al., 2024). As more populated cities tend to have larger surface areas as well, omitting to control for surface area underestimates the wage effect of population. Then, going from Column 2 to Column 4 by instrumenting for both density and surface area does lower the size of the density coefficient, because productivity $\tilde{A}_c t$ is an omitted variable that positively affects both city size and wages. Doubling

density, e.g., by moving from the average city of Clermont-Ferrand (196 inhab./km^2, c. 400,000 inhab.) to Nice (390 inhab./km^2, top decile of density, c. 1 million inhab.), entails an increase of about 3% in nominal wages.

Including instrumented $PM_{2.5}$ substantially changes the density coefficient (Column 6 against Column 4). As shown in Section 2.4, this change alone does not distinguish the amenity channel from the productivity channel, as both lower nominal wages. All else equal, cities with better air quality have higher productivity, which unambiguously implies higher wages. They also have better perceived amenity levels, and, because agglomeration economies are large enough, this also leads them to have higher nominal wages. These compounding effects imply that the air pollution generated by agglomeration erodes part of the very gains that agglomeration produces. At a given level of air quality, the wage gradient (0.044) is about 40% larger than the total gradient (0.031), and the difference is statistically significant. Consistent with the prior that $\beta^{\text{total}}_{N/S} = \beta_{N/S} + \beta_Z \, \phi_N$, this difference closely matches the product of the direct effect of $PM_{2.5}$ on wages and the elasticity of $PM_{2.5}$ to density estimated in Table 2: $-0.129 \times 0.115 \approx -0.015$, against $0.031 - 0.044 = -0.013$. This is comparable in magnitude to the role played by the spatial sorting of skilled workers in inflating the urban wage premium (about 30% according to Combes et al. (2008)).

The role of city surface area Columns 3–6 of Table 3 also shed light on the role of city surface area. When surface is not included, the density coefficient is smaller, since more populated cities also tend to be more spread out. More central to this paper, the coefficient on $PM_{2.5}$ is much larger in absolute value without surface (Column 5: -0.675) than with it (Column 6: -0.129). Since a larger surface lowers $PM_{2.5}$ at a given population, the specification without surface attributes to air pollution part of the wage effect of city extent. Finally, in the preferred specification of Column 6, the coefficient associated with city surface, $\hat{\beta}_S - \hat{\beta}_{N/S}$, is small and not statistically different from zero (-0.012, s.e. 0.011). This is in line with the theoretical model, in which agglomeration economies depend on population only, so that, conditional on air pollution, surface has no direct effect on wages.

Robustness checks Appendix Table D.3 shows a series of robustness checks of these reduced-form results by mirroring Table 3: these effects are similar looking at employment, i.e., city size N, instead of density N/S (Panel a), looking at city population density instead of employment density (Panel b), and looking at the full sample of cities, i.e., including those for which I cannot observe housing prices (Panel c). Appendix Table D.9 shows that these results are robust to using the commuting zone definition instead of the urban area. Finally, switching from a relative to an

absolute measure of $PM_{2.5}$ in Column 6 renders the exact same conclusions, as coefficients $\hat{\beta}_{N/S}$ and $\hat{\beta}_{S}$ are very similar to the main ones.

Reaction of working hours Table 3 focuses on adjusted hourly wages. This is indeed my object of interest, as it is the one that comes closest to a definition of labour productivity, and as I don't include working hour choices in the model (see Section 2.1 for a discussion on this point). Still, as my data allows, I provide a test for whether air pollution affects working hours in this cross-city long-run spatial equilibrium context. Appendix Table D.4 runs the same regressions, this time looking at raw working hours (Panel a), working hours adjusted for worker-level heterogeneity and differences in sectoral composition (Panel b), and labour earnings, adjusted for the same confounders (Panel c). The results first suggest that working hours are lower in denser cities, in line with higher hourly wage yielding an income effect.[40] The coefficient associated with density does not significantly vary once I control for air pollution in Column 6, and the coefficient for air pollution itself is insignificant. This is in line with absenteeism responses to density-generated air pollution being negligible in a cross-city spatial equilibrium context, as compared to labour productivity effects.

5.3 Air quality and the urban housing cost premium

Attenuation of the housing cost premium Table 4 shows the results for housing prices. Columns 1-2 show OLS results, and the rest of the columns show TSLS results. Coefficients associated with city size (resp., city surface) are consistently positive (resp., negative), in line with both the theoretical model and previous evidence on the costs of agglomeration (Combes et al., 2019). Column 4 shows that clean air capitalises into higher city-level housing prices, though not statistically significantly so. This is not a precise zero, which suggests some heterogeneity. The estimated elasticity (net of the impact of density and city surface), is around -0.22%, and is therefore close, though slightly smaller, to previous cross-city estimations for the U.S., obtained from temporal variation (Chay and Greenstone, 2005; Gruhl et al., 2025) or using hedonic valuation (Bayer et al., 2009). This is however larger than within-city estimates for Chicago (Chattopadhyay, 1999) or Tehran (Amini et al., 2022). In the French city of Nantes, Le Boennec and Salladarré (2017) find no evi-

[40]This partly goes against the evidence of Rosenthal and Strange (2008a) in the U.S., who find a negative reaction of working hours to density for low-skill workers, and a positive one for high-skill workers, for which returns to effort through labour market interactions, competition and career incentives are stronger. However, I condition on skill differences. More specifically, comparing Panel (a) to Panel (b) in Table D.4 reveals a more negative reaction of working hours to density once skill is controlled away, which suggests that the effort effect is larger for high-skill workers than for low-skill workers. In the French context, the nationwide 35-hour work week cap also leaves less room for such returns to effort through longer working hours.

Table 4: Local air pollution and the urban housing cost premium

	Log housing price index					
	(1)	(2)	(3)	(4)	(5)	(6)
(log) density	0.221***	0.263***	0.190**	0.377***	0.496***	0.446***
	(0.034)	(0.032)	(0.086)	(0.077)	(0.173)	(0.148)
(log) area		0.090***		0.061*		0.064*
		(0.018)		(0.033)		(0.033)
(log) $PM_{2.5}$					-1.34	-0.226
					(1.17)	(0.958)
Implied $\hat{\eta}_S = \hat{\beta}_S - \hat{\beta}_{N/S}$		-0.173		-0.316		-0.382
City size instrumented?			X	X	X	X
Area instrumented?			X	X		
$PM_{2.5}$ instrumented?				X		X
Year fixed effects	X	X	X	X	X	X
Observations	2,313	2,313	2,295	2,277	2,295	2,277
Clusters (cities)	257	257	255	253	255	253
Mean outcome	8.11	8.11	8.11	8.11	8.11	8.11
F-stat, (log) density			30.570	41.679	38.256	33.599
F-stat, (log) area				367.16		299.69
F-stat, (log) $PM_{2.5}$					33.995	37.434

Notes: Standard errors clustered at the city (MSA) level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$
Geography controls include mean elevation, ruggedness, and dummies for access to the sea, presence of a waterway, and mountains. Weather controls include summer and winter average temperature, wind speed, cloud cover and total precipitation. Geography controls include a dummy for access to the sea, to a waterway, to mountains, ruggedness and mean elevation. IV for density: soil characteristics. IV for area: 1806 population. IV for $PM_{2.5}$: number of days with an unusually low PBL height.

dence of such capitalisation. Obtaining larger estimates across cities than within cities is sensible considering that city-level air quality is often better perceived by households, who know that Paris is polluted, while Nantes or Bordeaux are "clean", while they likely perceive smaller within-city differences. Moreover, the cross-city variations estimated here reflect long-run, equilibrium differences in amenity levels, as studied in the theoretical framework, which likely generates higher elasticities than within-city variations.

Effect on real wages Appendix Table D.1 brings the results of Tables 3 and 4, by looking at the effect on real wages directly, defining the latter as nominal wages, minus the share of housing in expenditures $\beta = 0.33$, times the housing price index. First, putting air pollution aside,

I estimate a real wage "penalty" to city size: causal estimates in Column (4) show that a 1% increase in city population decreases real wages by about 0.044%. The literature in urban economics has long highlighted that the existence and growth of cities results from the balance of agglomeration and dispersion forces. (Net) agglomeration economies are thus theoretically bell-shaped: the centripetal force first prevails over the centrifugal one, but after a certain threshold, diseconomies and negative externalities arising from congestion and rising rents exceed the agglomeration economies (Fujita et al., 1999). This negative relationship between city size and real wages means that French cities are located on the right-hand side of the bell-shaped curve. This is despite a positive but small net effect of city surface (0.002), which slightly relaxes the surge in housing prices caused by higher density, but not sufficiently so. This is in line with the Rosen-Roback intuition of the theoretical framework, which states that lower amenity levels provided in smaller cities must be compensated for by higher real wages.

Columns (5)–(7) include instrumented $PM_{2.5}$. The real wage gradient at a given level of air quality (Column 6) is significantly less negative than the total gradient (Column 4): about 36% of the real wage penalty to density operates through the air pollution that density generates. This corresponds to a negative pollution component $\beta_Z^R \phi_N$, which indicates that the productivity effect of air pollution dominates its compensating effect. This holds despite the negative effect of air pollution on housing prices, which raises real wages through lower housing costs.

The role of city surface area When city surface is held fixed (Columns 4 and 6), about a third of the real wage penalty to density operates through the pollution channel ($1 - 28/44 = 36\%$). When city surface is allowed to vary, the pollution component accounts for less than 10% of the gradient ($1 - 39/42 = 7\%$). As the city can expand into space, a given increase in population generates a smaller increase in air pollution, and hence smaller wage losses. Table 4 shows the same pattern for housing prices: the housing price gradient at given air quality exceeds the total gradient by more when surface is held fixed (Column 6 against Column 4) than when it varies (Column 5 against Column 3). The implied coefficient on surface area is also more negative at a given level of air quality: when air pollution is allowed to vary, part of the dampening effect of sprawl on housing prices is offset by the improvement in air quality it brings, which raises housing prices.

Consistency check: Direct effect on (real) wages Appendix Table D.2 shows the effects of local air pollution on nominal (Columns 1-4) and real wages (Columns 5-7). The OLS estimates of Columns 1 and 2 are positive, as high local air pollution and high wages are both positively influenced by productivity-enhancing factors and city size. Turning to the IV results in Columns 3 and

4, the estimated elasticity of nominal wages to $PM_{2.5}$ lies around -0.16. In levels, this means that a city whose $PM_{2.5}$ concentration is one $\mu g/m^3$ higher offers nominal wages that are 1.7% lower. It is rather difficult to find comparable estimations in the literature. Borgschulte et al. (2024) examine the effect of an additional day of exposure to wildfire smoke on quarterly earnings in the US, and find a semi-elasticity of -0.1%.

Discussion A pure consumption-disamenity view predicts that real wages, measured as $W^{\text{adj}} - \beta\times$ (log) housing price index, should rise with pollution, compensating residents for the disamenity. The estimates of Columns 5 to 7 of Table D.2 do not support this, as they show a negative effect of $PM_{2.5}$ on real wages. As Section 2 shows, this can occur whenever the productivity cost of air pollution outweighs the amenity-compensation channel: a negative real-wage response is a sign of local air pollution acting as a production disamenity, not only a consumption one. However, as also highlighted in equations (18) and (20), these reduced-form results are not sufficient to be able to conclude on whether this is due to the wage capitalisation of *any* amenity or productivity shock being large (i.e., η_A^w and η_θ^w being large), or whether this is due to the effect on labour productivity being relatively large (i.e., ζ_A being large compared to ζ_θ).

6 Structural model: Calibration and fit

This section takes the spatial equilibrium model of Section 2 to the data. It starts by discussing how parameters that are shared with existing work are calibrated. Then, it combines the structure of the model to the reduced-form evidence of 5 in order to recover estimates of the effect of local air quality on perceived amenity levels, ζ_θ, and on labour productivity, ζ_a. All calibrated and estimated parameter values are summarised in Appendix Table A.5.

6.1 Calibration of common parameters

I calibrate, rather than estimate, the values of the model's deep parameters that do not pertain to the generation of local air pollution nor its effects. Indeed, the literature has recently produced reliable estimates of most of these, including in the French context.

First, worker-level parameters are the following. I set the share of expenditure dedicated to housing β to a cross-city constant of 0.33, and the cost of commuting τ to 0.15, with a negligible cost of within-CBD commuting of 0.05 (Combes et al., 2019). The commuting cost parameter also governs the within-city housing price gradient, which the model implies to be $q(x) \propto (b+x)^{-\tau/\beta}$, that is, an elasticity of $-\tau/\beta \approx -0.45$ with respect to distance. Estimating this gradient on DV3F

transactions located outside the central municipality yields an elasticity of -0.3 (Appendix Table B.2, Column 6), which implies $\tau \approx 0.10$, with a 95% confidence interval that includes 0.15. I keep $\tau = 0.15$: within this interval, it implies a weaker response of the city surface to housing prices (as discussed in Section 6.3 below), and is therefore conservative with respect to the land margin. Then, noting that the way city population enters the labour supply equation shown in (4) is isomorphic to that of Ahlfeldt et al. (2024), I set the congestion parameter δ to $1/\epsilon$, where ϵ corresponds to the dispersion of idiosyncratic tastes for a city over another. The closer ϵ is to 0, the more one worker prefers one city over another, and the less elastic the labour supply. As ϵ tends to $+\infty$, we have $\delta = 1/\epsilon = 0$ and labour is perfectly elastic, which goes back to the classical Rosen-Roback case without spatial frictions. At the cross-city level, ϵ is measured to be between 2 and 3 (Monte et al., 2018; Bilal, 2023). I thus take $\delta = 1/\epsilon = 1/2.5 = 0.4$.

Second, housing production depends on available land (extensive margin) and the cost of building taller, σ (intensive margin). As available land is determined through the value of the agricultural land rent, the housing equilibrium equation (7) only depends on the building cost, with $\kappa = (\sigma - 1)/\sigma$. I thus calibrate κ following work on intensive housing supply elasticities, and set it to $\kappa = 0.3$ (Combes et al., 2021; Baum-Snow and Duranton, 2025).

Third, turning to firms, I set the labour share in production α to 0.75 and the capital share γ to 0.2, from Cette et al. (2019) and Gutiérrez and Piton (2020). The agglomeration parameter is not directly estimated, but I don't borrow it from the literature either. I deduce its value from the structural equation (16), which implies that $\lambda = \frac{(1+\hat{\beta}_N)(1-\gamma)}{\alpha} - 1$. I use my estimate of $\hat{\beta}_N$ from Column (4) in Table 5, and the calibrated values of α and γ from above to obtain $\hat{\lambda} = 0.11$.

6.2 Production *vs* consumption disamenity: Deducing ζ_θ and ζ_a

Second, I need to estimate parameter values that govern the endogenous formation of local air pollution and its feedback onto local fundamentals. On the one hand, the reaction of pollution Z to population and surface area, i.e. parameters ϕ_N and ϕ_S, are directly estimated in the reduced-form results of Table 2. On the other hand, the effects of Z on amenities θ and labour productivity $\tilde{A}$ are inferred by applying the model's structure onto the reduced-form coefficients.

Deducing (dis)amenity values Given the structure of the model, I am able to deduce the value of ζ_θ and ζ_a from the estimates of Table 3. I rewrite equations (18) and (33) into a system, and

equalise them to the reduced-form estimates of Section 5:

$$\begin{cases} \frac{d\ln(w^*_{\tilde{S},ct})}{d\ln(N^*_{\tilde{S},ct})} & = \frac{d\ln(w^*_{\tilde{S},ct})/d\ln(\theta_{ct})}{d\ln(N^*_{\tilde{S},ct})/d\ln(\theta_{ct})} = \frac{\eta^w_\theta - \left(\eta^w_a z_a + \eta^w_\theta z_\theta\right)\phi_N \ \eta^N_\theta}{\eta^N_\theta\left(1-\left(\eta^N_a z_a + \eta^N_\theta z_\theta\right)\phi_N\right)} = \hat{\beta}^w_{N|S} \\ \frac{d\ln(q(0)^*_{\tilde{S},ct})}{d\ln(N^*_{\tilde{S},ct})} & = \frac{d\ln(q(0)^*_{\tilde{S},ct})/d\ln(\theta_{ct})}{d\ln(N^*_{\tilde{S},ct})/d\ln(\theta_{ct})} = \frac{\eta^q_\theta - \left(\eta^q_a z_a + \eta^q_\theta z_\theta\right)\phi_N \ \eta^N_\theta}{\eta^N_\theta\left(1-\left(\eta^N_a z_a + \eta^N_\theta z_\theta\right)\phi_N\right)} = \hat{\beta}^q_{N|S} \end{cases} \tag{30}$$

I rely on the calibration of Table A.5 discussed above to derive the values of the η coefficients up to the values of the ζ parameters. This way, I am then able to derive the underlying values of ζ_A and ζ_θ from the combination of known parameters and the total gradients of wages ($\hat{\beta}^w_{N|S}$, Column 3 of Table 3) and housing prices ($\hat{\beta}^q_{N|S}$, Column 3 of Table 4), estimated in the specifications without $PM_{2.5}$. I deduce from these that $\zeta_A = 0.11$ and $\zeta_\theta = 0.06$.

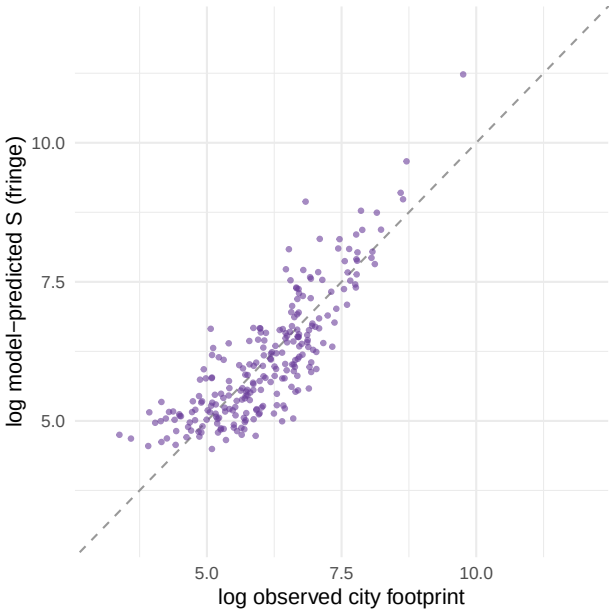
As a consistency check, the total gradients estimated without $PM_{2.5}$ and without surface area can also be matched to the equations, shown in the Appendix, of the urban wage and housing cost premia with an endogenous city surface. Following this route, I find consistent estimates of ζ_A ad ζ_θ. Those lie at 0.11 for the former, and 0.03 for the latter.

What do ζ_θ and ζ_a capture? I argue that this estimation method allows me to isolate the effect of city structure-generated differences in air pollution levels on perceived amenities and productivity across cities. This relies on the assumption that any other city characteristic that is dependent on city structure and affects local amenities and/or productivity is filtered out.

With respect to ζ_a, one could think of three potentially confounding elements, that are all taken care of by the identification strategy. First, industrial composition and worker sorting may be at play, with cities specialised in dirtier sectors having both higher air pollution and wage levels. This is accounted for through the two-step method used to construct the wage measure, which nets worker and industry effects out of the wage premium before it feeds the calibration. Second, there could be other production amenities orthogonal to air pollution that provide a natural advantage to denser or sprawling cities. This is absorbed by the λ parameter, which is calibrated to match my observed wage premium that controls for air pollution. Third, and finally, unobserved pollutants like NO_2 and O_3 could be generated by density and then have an effect on amenities and productivity. These co-pollutants share a common generator with $PM_{2.5}$, car traffic, which is itself a function of city structure. Like in the reduced-form results, ζ_a could therefore be argued to be a productivity effect of this pollution bundle, and I do not claim to separate them out.

With respect to ζ_θ, one could think of several confounding elements, one of which is also co-pollutants, for which the same reasoning as for ζ_a applies. The first specific component is traffic noise, which is a function of city size, and hits the amenity margin. The estimated ζ_θ here is

Figure 2: Model-predicted city surface area against observed city surface area



Notes: One dot for each of 248 cities. The 45-degree line is dotted. Observed city surface area from INSEE 2010 MSA definition. Model-predicted city surface area from main calibration according to Table A.5. The slope is equal to 0.882, the R^2 to 0.693, and the coefficient of correlation to 0.832.

stripped out of this confounder through controlling for the commuting cost term $\tau(x)$. The second potential threat relates to consumer-city amenities, which rise with city size and capitalise into housing prices (Glaeser et al., 2001). This is cancelled out under the reasonable assumption that, conditional on city size and surface, consumer-city amenities are orthogonal to air quality. The third potential element is crime, which tends to rise with city size (Glaeser et al., 1996; Glaeser and Sacerdote, 1999), although a causal link has recently been disputed (Exbrayat and Stephane, 2025). Again, this does not affect my estimate of ζ_θ under the assumption that crime is orthogonal to air pollution levels, conditional on density.[41]

6.3 Assessing model fit

As the model gives expressions for their equilibrium values, I solve the model by exploiting data on observed city population, wages, housing prices and local air pollution. As such, by construction, these elements are well predicted by the model. However, in the version of the model where the city fringe is endogenised, I back the latter out using the same observed information. In other

[41] To some extent, it would be relevant for part of the crime component to be captured in ζ_θ, since air pollution was shown to have a positive effect on crime rates (Bondy et al., 2020; Herrnstadt et al., 2021). In that sense, stripping out differences in crime rates by introducing such a (likely bad) control would actually be removing part of the actual effect of air pollution on housing prices, and thus, perceived amenity levels.

words, city surface is pinned by the fundamentals using equation (24), such that its predicted values provide a test of the model's fit. I check the accuracy of predicted city surface by plotting it against observed city surface area in Figure 2. Overall, the model matches observed surface areas well, although with a slight over-prediction at the tails of the distribution.

Figure 2 assesses the level of city surface, not its response to shocks, which is what matters for the counterfactuals. In the model, the fringe responds only to housing prices at the CBD: given construction productivity and agricultural rent, $d\ln S \approx (2\beta/\tau)\, d\ln q(0)$, an elasticity of about 4.4 under the calibration. This response abstracts from land use regulation. Restrictions on the conversion of agricultural and natural land, as the no net land take objective, create a wedge at the fringe: they lower the effective elasticity without affecting the within-city price gradient. City surface responses in the endogenous-surface model should therefore be read as an upper bound, and the fixed-surface model as the closed-margin bound.

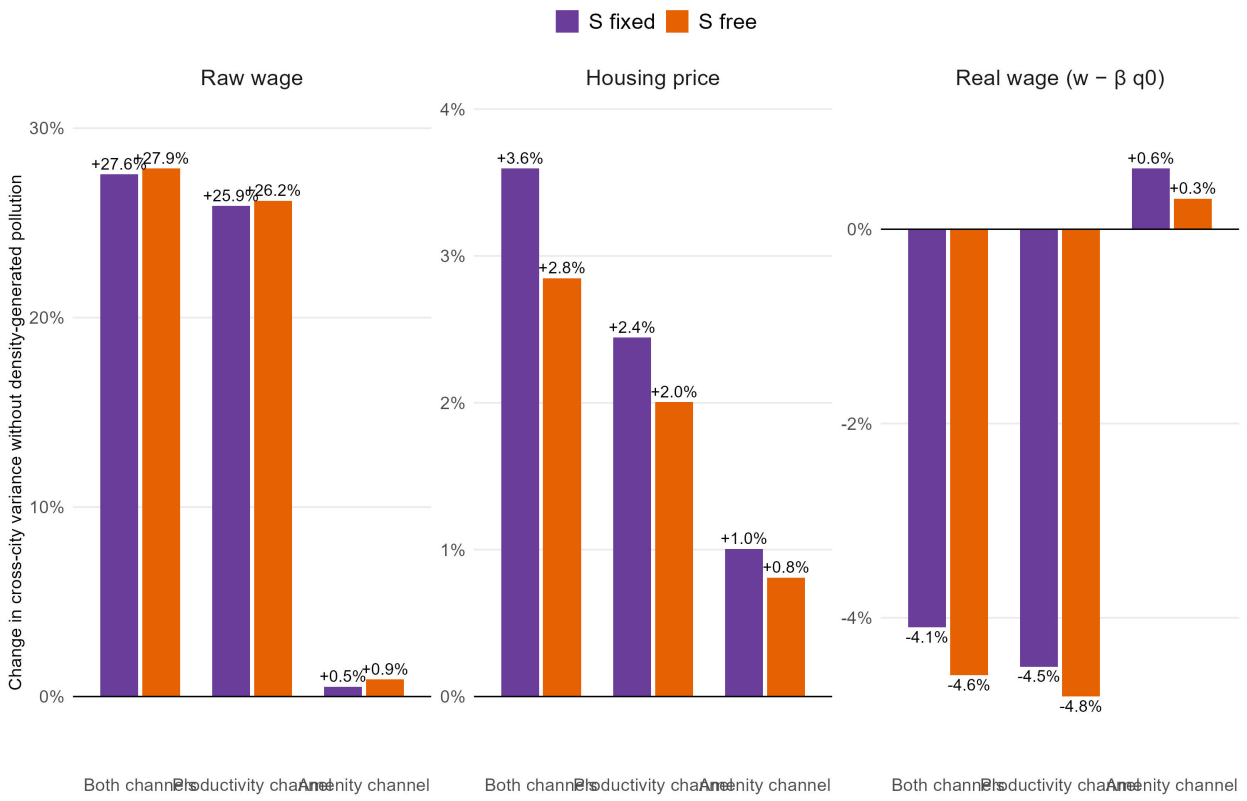
7 Counterfactual exercises

I provide estimations of the effect of density-generated air pollution on spatial wage disparities, as well as counterfactual exercises based on credible scenarios of city-specific population growth and local pollution responsiveness to a population increase. For each counterfactual simulation, I look at the effect both holding city surface constant, and allowing it to move, so as to shed light on the role of sprawl as a mitigator of the air pollution channel.

7.1 Air pollution and spatial (real) wage disparities

The model built in Section 2 allows for a general-equilibrium estimation of the role played by local air quality in shaping wage, housing price, and real wage differentials across cities. Figure 3 shows the results of a counterfactual exercise taking the variance of observed raw wages (top left), wages adjusted for worker sorting (top right), housing prices (bottom left) and real wages (bottom right). Starting from the estimated baseline, I sever the link between density and air pollution ($\phi_N = 0$), while holding the damages caused by air pollution fixed. I then decompose this effect by letting the reduction in air pollution operate through productivity only, and through amenities only. Each bar shows the share of baseline cross-city variance that is removed, computed as one minus the ratio of counterfactual to baseline variance. I do this for a fixed city surface in purple, and for an endogenous one in orange. A negative value means that air pollution from city density compresses dispersion, so that tackling it widens the distribution; a positive value means the reverse.

Figure 3: Which fraction of the variance in wages and housing prices is attributable to air pollution?



Notes: Results from a counterfactual exercise shutting local air pollution down (left), or only the feedback on labour productivity (middle), or only the feedback on amenity level (right). Figures shown are the change in the variance in wages (top) and housing prices (bottom) across the 248 largest French cities.

Air pollution damages from density account for a large share of the cross-city dispersion in wages. Shutting the generation link down entirely removes 27.6% of the baseline variance in raw wages and 34.6% in adjusted wages under a fixed surface (27.9% and 35.0% under a free one). The negative sign is central: density-generated air pollution compresses the cross-city wage distribution, because more air pollution is generated in the densest cities. This drags down their amenities and productivity, which in turn pulls their wages towards those of smaller cities. The same holds, to a much smaller extent, for housing prices (3.6% of baseline variance). Air pollution thus acts as a spatial equaliser across the city size distribution. If a policy allowed cities to generate less pollution from their density, for instance through cleaner transport, electrified heating or tighter emission controls, this force would weaken and the largest cities would pull away again.

Real wages go the other way: the share of variance explained by air pollution is positive, meaning that abating air pollution from density compresses real wage dispersion rather than widens

it. The reason is that the compression of nominal wages and of housing prices do not cancel once housing costs are netted out: in real terms, density-generated pollution widens the income distribution across cities. Severing the generation link narrows real wage dispersion by 4.1% under a fixed surface, and by 4.6% under a free one. Air quality is thus equalising in nominal terms but disequalising in real, cost-of-living terms.

The productivity channel accounts for almost all of these effects: on its own, it removes 25.9% of the variance in raw wages, 32.6% in adjusted wages, and narrows real wage dispersion by 4.5% (fixed surface). The amenity channel plays a minor role, with effects below 1.2% in absolute value for all outcomes. This asymmetry stems only partly from the damage elasticities, as ζ_a is less than twice as large as ζ_θ (0.11 vs. 0.06), and mostly from how each fundamental passes through to wages. A productivity shock shifts labour demand and passes almost one-for-one into wages ($\eta_A^w = 1.01$). An amenity shock, by contrast, only affects wages through the inflow of workers it attracts. This inflow has two opposite effects on wages: a higher labour supply lowers wages through diminishing returns to labour, while agglomeration economies raise productivity.[42] By contrast, both fundamentals make a city more attractive and are capitalised into housing prices with elasticities of similar magnitude ($\eta_A^q = 0.85$ and $\eta_\theta^q = 0.68$), so that the two channels contribute more evenly to the compression of housing prices.

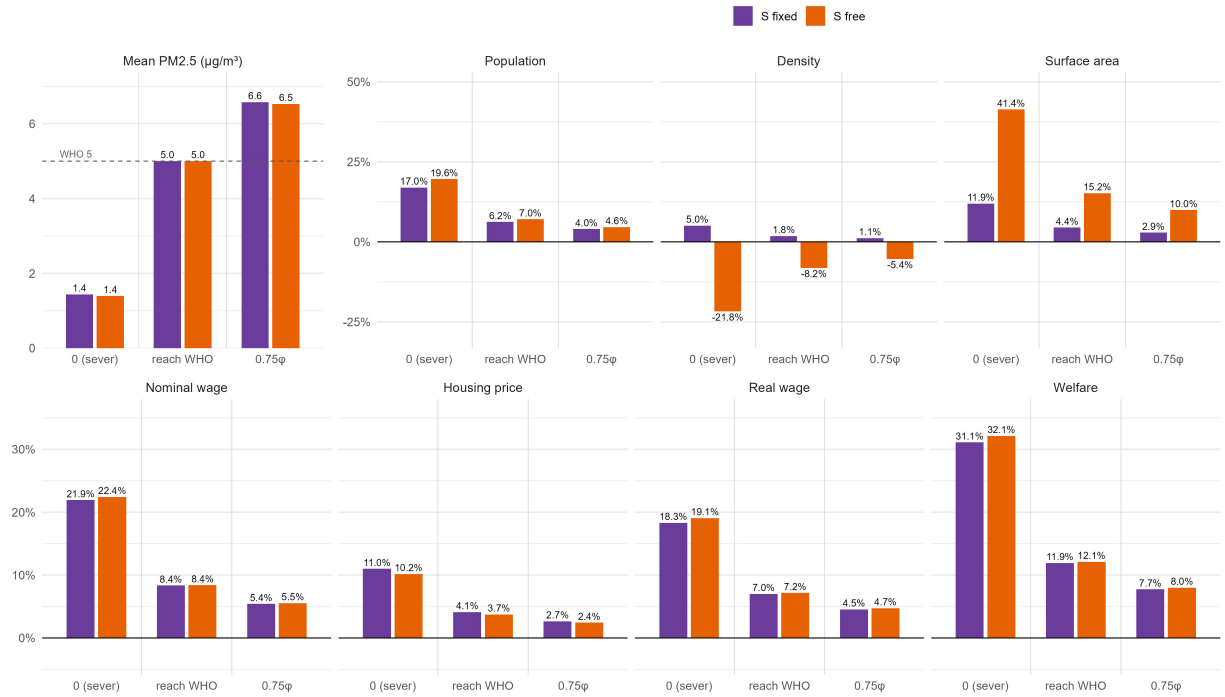
Finally, the fixed-surface bars should be read as the constrained, shorter-run response, and the free-surface bars as the longer-run response once spatial extent re-optimises. Allowing for sprawl leaves the wage results almost unchanged, but dampens the compression of housing prices (2.8% vs. 3.6%) and slightly strengthens that of real wages (4.6% vs. 4.1%): when a city can resize its spatial extent, it absorbs part of a population change by spreading out rather than through housing prices.

7.2 Bringing cities within the WHO guideline

In 2021, the WHO revised the recommended guideline for the average yearly concentration of fine particulates down from 10 $\mu g/m^3$ to 5 $\mu g/m^3$. In my sample in urban France, 2018 concentration is 9.85 $\mu g/m^3$ at the city level, and 10.94 at the worker level, i.e., around the 2005 guideline of 10 $\mu g/m^3$. Figure 4 shows three counterfactual exercises, or scenarios, around this objective to quantify whether a reduction of the elasticity of local $PM_{2.5}$ to density would allow to reach it. It shows the reaction of raw wages, adjusted wages, housing prices, real wages, population, density, city surface, and, of course, $PM_{2.5}$. All results are displayed at the worker level, by weighting each

[42] With my calibrated parameters, the two nearly offset each other ($\alpha\lambda = 0.08$ against $1 - \alpha - \gamma = 0.05$): a 1% amenity-driven increase in population raises wages by only 0.04% (equation **??**), hence the small $\eta_\theta^w = 0.09$.

Figure 4: Meeting the WHO guideline by reducing the elasticity of air pollution to density



Notes: Counterfactual exercises whose target is to set the overall mean level of population-weighted $PM_{2.5}$ below the 5 µg/m^3 WHO guideline. Results shown using city-level population weights at the new equilibrium, such that they are at the worker level, for the average at the new equilibrium. Current-day (2018) $PM_{2.5}$ is equal to 10.94 µg/m^3 at the worker level.
Scenario 1 turns off the city structure-air pollution channel entirely, to set a benchmark, by taking the observed $\hat{\phi}_N$ down from 0.15 to 0. Scenario 2 targets population-weighted mean $PM_{2.5}$ to be equal to the 5 µg/m^3 WHO guideline, and deduces that under a fixed (resp., free) city surface, the observed $\hat{\phi}_N$ should reach 0.092 (resp., 0.093). Scenario 3 reduces the observed $\hat{\phi}_N$ by a quarter, from 0.15 to 0.11.

city by its new equilibrium population, as this is most policy-relevant.

The first scenario severs the link between urban density and local $PM_{2.5}$, by setting $\phi_N = 0$.[43] For each outcome, this corresponds to the two bars on the left-hand side. This scenario would meet the WHO guideline, as it achieves a mean $PM_{2.5}$ of 1.4 µg/m^3 under both a fixed and a free city surface. Therefore, it generates average wage gains of 21.7 to 22.4%, and real wage gains of 18.3 to 19.1%, since the reaction of housing prices is more limited (10.2 to 11.0%). As air pollution decreases in the most dense, polluted and populated cities, new workers are drawn there, such that population and density increase in the already largest cities. The latter also expand more into space.[44]

As a perfect decoupling of density and air pollution is unrealistic, I use the scenario where

[43] It matches the scenario in Section 7.1, this time looking at mean worker outcomes as opposed to city variances.

[44] At the worker level, average city surface would also increase in the scenario where it is fixed, because of a displacement of workers from more compact cities to more sprawling ones.

$\phi_N = 0$ as a benchmark for two other exercises. The second scenario takes 5 $\mu g/m^3$ as a target for mean $PM_{2.5}$, and backs out the elasticity of $PM_{2.5}$ to density needed to reach it: ϕ_N must fall by 38% to 0.092. Reaching the guideline would raise wages by about 8% and real wages by 7.0 to 7.2%. Notice that real wage gains are slightly larger under a free surface, as cities accommodate the inflow of workers by sprawling (+15.2% in surface area) rather than densifying, which dampens the reaction of housing prices (3.7% vs. 4.1%).

The third scenario reduces the observed elasticity by a quarter, to $\phi_N = 0.11$. Mean $PM_{2.5}$ then reaches 6.6 $\mu g/m^3$, above the WHO guideline, but still yields wage gains of 5.4 to 5.5% and real wage gains of 4.5 to 4.7%, i.e., about a quarter of the benchmark gains. Overall, reducing the elasticity of air pollution to density generates substantial gains.

Appendix Figure E.1 shows the same results as Figure 4, but at the city level in Panel (a), and weighting by baseline city population instead of new-equilibrium city population in Panel (b). Appendix Figure E.2 decomposes these effects into the productivity and amenity channels at the worker level, and E.3 at the city level. The productivity channel accounts for virtually all of the gains in nominal and real wages, while the amenity channel accounts for a substantial share of the reallocation of workers (about 40% of the population increase under a fixed surface when $\phi_N = 0$) and part of the increase in housing prices.

7.3 Projecting current internal migration trends

In a third exercise, I back out the change in local amenities that explains the observed population trend from 1999 to my last year of observation, 2018. I then feed these changes in perceived amenities back into the model to predict the effect that a continuation of the current migration trend from cities of the North and East to the Atlantic and Mediterranean coasts, would have on wages and housing prices, and the share of these effects attributable to induced changes in local air pollution.[45] Panel (E.4a) of Appendix Figure E.4 shows the observed 1999-2018 changes in population, and Panel (E.4b) the amenity changes that the model requires to explain them. Appendix E.2 provides more details on the technical steps.

Figure 5 shows the results of this exercise at the worker level (i.e., weighting by new equilibrium city population).[46] Under today's elasticities, continuing the 1999-2018 trend barely affects

[45] The model, and thus the counterfactual exercise, abstract from modelling in-migration from outside cities, and treats them as a closed system. In the projection exercise, each city's target population embeds its observed 1999-2018 trend, including in-migration from rural areas and from abroad, such that total urban population grows as observed. The model does not, however, specify where new residents come from, nor allow air pollution to affect the total inflow into the urban system, which would require an outside option

[46] Appendix Figure E.5 shows the results at the city level (Panel E.5a) and weighted by baseline city population (Panel E.5b).

Figure 5: Projecting current internal migration trends



Notes: Projection of 1999-2018 city population trends from 2018 on. Step 1: amenity shocks are backed out to reproduce each city's projected population. They average to zero across cities. Air pollution is held at its 2018 level, so that the shocks are identical across scenarios. Step 2 (shown in the figure): the model is solved under these shocks, with air pollution responding to the change in population and surface area, at today's pollution elasticities (ϕ) and with both ϕ_N and ϕ_S reduced by a quarter (0.75ϕ). Purple: fixed city surface area. Orange: endogenous city surface area. Bars show changes relative to the 2018 equilibrium, averaged across 248 cities and weighted by new projected equilibrium population.

average exposure: mean $PM_{2.5}$ reaches 11.1 $\mu g/m^3$ under a fixed surface area and 10.9 under a free one, against 10.94 in 2018. Even though density increases by 13.4% at the worker level when cities cannot expand into space, the elasticity of $PM_{2.5}$ to density is moderate enough for this to translate into an increase of only about 1.5% in concentration. The current internal migration trend is therefore neither a threat to, nor a path towards, the WHO guideline.

Nominal wages barely move under this trend (+0.4 to +0.7%), while housing prices rise by 4.9 to 5.8%, so that real wages fall by 0.9 to 1.5%. This is the counterpart, at the scale of the urban system, of the asymmetry discussed in Section 7.1: the trend is rationalised by amenity shocks, which pass through to wages only weakly but are capitalised into housing prices. Workers moving towards the Atlantic and Mediterranean coasts thus pay for the amenities of these cities through higher housing costs rather than lower nominal wages. The fall in real wages should therefore not be read as a welfare loss, but as the compensating differential for the amenities that attract

workers to these cities.

Whether city surface area can adjust determines how this growth is absorbed. When it is free, the trend is absorbed almost entirely through sprawl: average surface area increases by 17.0% at the worker level, of which about 3 percentage points reflect the reallocation of workers towards cities with larger surface areas (as shown by the fixed-surface bar), and average density does not change (-0.3%). When it is fixed, which is closer to the no net land take objective (ZAN), the same growth is absorbed through a 13.4% increase in density. Absorbing this growth through density rather than sprawl raises mean exposure by only 0.2 $\mu g/m^3$ of $PM_{2.5}$ (about 2%), which is small compared with the 6 $\mu g/m^3$ gap to the WHO guideline. Real wages fall by only 0.6 percentage point more under a fixed surface (-1.5 vs. -0.9%), mostly because housing prices rise more (5.8 vs. 4.9%), rather than because of air pollution.

Reducing both elasticities of air pollution to city structure by a quarter reverses the sign of the real wage change, to +3.1% under a fixed surface and +3.8% under a free one. Relative to the trend under today's elasticities, this partial decoupling raises real wages by 4.6 to 4.7 percentage points and brings mean $PM_{2.5}$ down to 6.5 to 6.7 $\mu g/m^3$. This is very close to the effect of a similar reduction starting from the 2018 equilibrium (real wage gains of 4.5 to 4.7%, Section 7.2). The gains from decoupling therefore barely depend on where population goes. They also operate through a different channel than the migration trend: decoupling mostly raises labour productivity, and hence nominal wages (+5.8 to 6.3%), whereas the trend mostly redistributes amenities and housing costs across cities.

Decoupling, however, adds to the pressure on housing in growing cities. With a fixed surface, housing prices rise by 8.5%, against 7.3% when cities can sprawl, as cleaner air makes cities more attractive and the inflow of workers can only be accommodated through higher prices. Combining cleaner air with a binding constraint on land take thus makes housing costs the main margin of adjustment, such that policies cutting emissions per resident would need to be paired with policies that expand housing supply within the existing surface.

Appendix Figure maps these changes city by city under a fixed surface, which approximates the no net land take objective. Population, and hence density, grows most along the Atlantic arc (Rennes, Nantes, Bordeaux), in Toulouse, Montpellier and the French side of the Geneva area, while cities of the North-East and the Centre shrink. As it is generated by density, local air pollution follows the same pattern, with increases of at most 4%. As nominal wages are virtually unchanged everywhere, the spatial pattern of real wages mirrors that of housing prices, which rise by over 7% in the fastest-growing cities: real wages fall most, by about 2%, in Toulouse, Montpellier and close to Switzerland, and barely move in declining cities.

8 Conclusion

This paper studies how the local air pollution that city structure generates shapes wage, housing price, and real wage differences across French cities. Its central finding is that density is partly self-taxing: the local air pollution it generates acts primarily as a production disamenity, which erodes part of the productivity gains that agglomeration delivers.

First, I provide reduced-form evidence that $PM_{2.5}$ rises with density, while a larger spatial extent lowers it by about as much at a given population. At a given level of air quality, the urban wage premium is about 40% larger than when air quality adjusts with density, and the air pollution that density generates accounts for about a third of the real wage penalty to density. Real wages fall with air pollution, such that the productivity cost of air pollution outweighs the compensation that workers receive for it. Second, I develop and estimate a spatial equilibrium model with endogenous air quality, which confirms that the productivity channel dominates. Counterfactual simulations show that density-generated air pollution compresses cross-city differences in nominal wages, but widens differences in real wages. Bringing French cities within the WHO guideline by weakening the link between density and air pollution alone would require a near-complete decoupling.

These results do not argue against densification, which remains justified on climate and productivity grounds, but they show that its local environmental cost is borne by the workers of the densest cities. Densification should therefore be paired with policies that cut emissions per resident, such as cleaner residential heating and public transport scaled to population growth, and with policies that reduce the level of emissions. Sprawl offers only a partial alternative: it dilutes air pollution but moves workers further from jobs, and with restrictions on land take such as the European "no net land take" objective, densification is likely to be the main margin of city growth in the coming decades. The magnitudes found here also likely represent a lower bound for rapidly urbanising settings, where the elasticity of air pollution to density is likely higher.

Finally, the model abstracts from worker heterogeneity, such that welfare is measured through a common utility level. Introducing skill differences would allow to quantify the distributional effects of density-generated air pollution across skill groups. Since high-skilled workers sort into the largest and most polluted cities, abating air pollution there would likely strengthen this sorting, with ambiguous effects on spatial inequality in real wages. Other dimensions of urban morphology, such as the distribution of density within cities, would also deserve attention from future research.

References

Adhvaryu, A., N. Kala, and A. Nyshadham (2022). Management and shocks to worker productivity. *Journal of Political Economy 130*(1), 1–47.

Ahlfeldt, G., F. Bald, D. Roth, and T. Seidel (2024). Measuring quality of life under spatial frictions. IZA Discussion Paper No. 17549.

Ahlfeldt, G. M. and E. Pietrostefani (2019). The economic effects of density: A synthesis. *Journal of Urban Economics 111*, 93–107.

Ahlfeldt, G. M., S. J. Redding, D. M. Sturm, and N. Wolf (2015). The economics of density: Evidence from the Berlin Wall. *Econometrica 83*(6), 2127–2189.

Ai, H., X. Zhang, and Z. Zhou (2023). The impact of greenspace on air pollution: Empirical evidence from China. *Ecological Indicators 146*, 109881.

Albouy, D. (2016). What are cities worth? Land rents, local productivity, and the total value of amenities. *Review of Economics and Statistics 98*(3), 477–487.

Albouy, D., W. Graf, R. Kellogg, and H. Wolff (2016). Climate amenities, climate change, and American quality of life. *Journal of the Association of Environmental and Resource Economists 3*(1), 205–246.

Alonso, W. (1964). *Location and Land Use: Toward a general theory of land rent*. Harvard University Press.

Amini, A., K. Nafari, and R. Singh (2022). Effect of air pollution on house prices: Evidence from sanctions on Iran. *Regional Science and Urban Economics 93*, 103720.

Basaglia, P. and L. Sarmiento (2025). When the Boundary Layer Drops: Air Quality and Healthcare Use in Mexico. CESifo Working Paper No. 11901.

Baum-Snow, N. and G. Duranton (2025). Housing supply and housing affordability. In *Handbook of Regional and Urban Economics*, Volume 6. Elsevier.

Bayer, P., N. Keohane, and C. Timmins (2009). Migration and hedonic valuation: The case of air quality. *Journal of Environmental Economics and Management 58*(1), 1–14.

Berson, C., P.-P. Combes, L. Gobillon, and A. Sotura (2024). Time-varying agglomeration economies and aggregate wage growth. Banque de France Working Paper No. 939.

Bieri, D. S., N. V. Kuminoff, and J. C. Pope (2023). National expenditures on local amenities. *Journal of Environmental Economics and Management 117*, 102717.

Bilal, A. (2023). The Geography of Unemployment. *Quarterly Journal of Economics 138*(3), 1507–1576.

Blaudin de Thé, C., B. Carantino, and M. Lafourcade (2021). The Carbon 'Carprint' of Urbanization: New Evidence from French Cities. *Regional Science and Urban Economics 89*, 103693.

Bondy, M., S. Roth, and L. Sager (2020). Crime is in the air: The contemporaneous relationship between air pollution and crime. *Journal of the Association of Environmental and Resource Economists 7*(3), 555–585.

Borck, R. and P. Schrauth (2021). Population density and urban air quality. *Regional Science and Urban Economics 86*, 103596.

Borgschulte, M., D. Molitor, and E. Y. Zou (2024). Air pollution and the labor market: Evidence from wildfire smoke. *Review of Economics and Statistics*, 1–46.

Brehm, J., N. Pestel, S. Schaffner, and L. Schmitz (2025). From Low Emission Zone to academic track: Environmental policy effects on educational achievement in elementary school. *Journal of Environmental Economics and Management 132*, 103165.

Burnett, R., H. Chen, M. Szyszkowicz, N. Fann, B. Hubbell, C. A. Pope III, J. S. Apte, M. Brauer, A. Cohen, S. Weichenthal, et al. (2018). Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter. *Proceedings of the National Academy of Sciences 115*(38), 9592–9597.

Carneiro, J., M. A. Cole, and E. Strobl (2021). The effects of air pollution on students' cognitive performance: Evidence from Brazilian university entrance tests. *Journal of the Association of Environmental and Resource Economists 8*(6), 1051–1077.

Carozzi, F. and S. Roth (2023). Dirty density: Air quality and the density of American cities. *Journal of Environmental Economics and Management 118*, 102767.

Cette, G., L. Koehl, and T. Philippon (2019). The labor share in the long term: A decline? *Economics and Statistics* (510-511-512), 35–51.

Chang, T. Y., J. Graff Zivin, T. Gross, and M. Neidell (2016). Particulate pollution and the productivity of pear packers. *American Economic Journal: Economic Policy 8*(3), 141–69.

Chang, T. Y., J. Graff Zivin, T. Gross, and M. Neidell (2019). The effect of pollution on worker productivity: Evidence from call center workers in China. *American Economic Journal: Applied Economics 11*(1), 151–72.

Chattopadhyay, S. (1999). Estimating the demand for air quality: New evidence based on the Chicago housing market. *Land Economics*, 22–38.

Chay, K. Y. and M. Greenstone (2003). The impact of air pollution on infant mortality: Evidence from geographic variation in pollution shocks induced by a recession. *Quarterly Journal of Economics 118*(3), 1121–1167.

Chay, K. Y. and M. Greenstone (2005). Does air quality matter? Evidence from the housing market. *Journal of Political Economy 113*(2), 376–424.

Chen, Z., P.-P. Combes, S. Démurger, and X. Liu (2024). Housing Costs and Real Income Differences across Chinese Cities. SciencesPo Economics Discussion Paper No. 2023-11.

Ciccone, A. and R. E. Hall (1996). Productivity and the Density of Economic Activity. *American Economic Review 86*(1).

Citepa (2018). Gaz à effet de serre et polluants atmosphériques: Bilan des émissions en France de 1990 à 2017.

Combes, P.-P., G. Duranton, and L. Gobillon (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics 63*(2), 723–742.

Combes, P.-P., G. Duranton, and L. Gobillon (2012). The Costs of Agglomeration: Land Prices in French Cities. CEPR Discussion Paper No. 9240.

Combes, P.-P., G. Duranton, and L. Gobillon (2019). The Costs of Agglomeration: House and Land Prices in French Cities. *Review of Economic Studies 86*(4), 1556–1589.

Combes, P.-P., G. Duranton, and L. Gobillon (2021). The production function for housing: Evidence from France. *Journal of Political Economy 129*(10), 2766–2816.

Combes, P.-P., G. Duranton, L. Gobillon, C. Gorin, and F. Robert-Nicoud (2024). Urbanisation and urban divergence: France 1760-2020.

Combes, P.-P., G. Duranton, L. Gobillon, and S. Roux (2010). Estimating Agglomeration Economies with History, Geology, and Worker Effects. In *Agglomeration Economics*, pp. 15–66. University of Chicago Press.

Conwell, L. J., F. Eckert, and A. M. Mobarak (2023). More roads or public transit? Insights from measuring city-center accessibility. NBER Working Paper No. 30877.

Copeland, B. R. and M. S. Taylor (2004). Trade, Growth, and the Environment. *Journal of Economic Literature 42*(1), 7–71.

Deryugina, T., G. Heutel, N. H. Miller, D. Molitor, and J. Reif (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review 109*(12), 4178–4219.

Diamond, R. (2016). The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *American Economic Review 106*(3), 479–524.

Diamond, R. and E. Moretti (2026). Where is standard of living the highest? local prices and the geography of consumption. Working Paper.

Duranton, G. and D. Puga (2004). Micro-foundations of urban agglomeration economies. In *Handbook of Regional and Urban Economics*, Volume 4, pp. 2063–2117. Elsevier.

Ebenstein, A., V. Lavy, and S. Roth (2016). The long-run economic consequences of high-stakes examinations: Evidence from transitory variation in pollution. *American Economic Journal: Applied Economics 8*(4), 36–65.

Exbrayat, N. and V. Stephane (2025). Does urbanization cause crime? Evidence from rural–urban migration in South Africa. *Journal of Urban Economics 149*, 103787.

Fowlie, M., E. Rubin, and R. Walker (2019). Bringing satellite-based air quality estimates down to earth. *AEA Papers and Proceedings 109*, 283–288.

Fujita, M., P. Krugman, and A. Venables (1999). *The Spatial Economy: Cities, regions, and international trade*. MIT Press.

Ganong, P. and D. Shoag (2017). Why has regional income convergence in the US declined? *Journal of Urban Economics 102*, 76–90.

Glaeser, E. L. (2010). *Agglomeration Economics*. University of Chicago Press.

Glaeser, E. L. and J. D. Gottlieb (2009). The wealth of cities: Agglomeration economies and spatial equilibrium in the United States. *Journal of Economic Literature 47*(4), 983–1028.

Glaeser, E. L., J. Kolko, and A. Saiz (2001). Consumer city. *Journal of Economic Geography 1*(1), 27–50.

Glaeser, E. L. and D. C. Mare (2001). Cities and skills. *Journal of Labor Economics 19*(2), 316–342.

Glaeser, E. L. and B. Sacerdote (1999). Why is there more crime in cities? *Journal of Political Economy 107*(S6), S225–S258.

Glaeser, E. L., B. Sacerdote, and J. A. Scheinkman (1996). Crime and social interactions. *Quarterly Journal of Economics 111*(2), 507–548.

Godzinski, A. and M. Suarez Castillo (2021). Disentangling the effects of air pollutants with many instruments. *Journal of Environmental Economics and Management 109*, 102489.

Graff Zivin, J. and M. Neidell (2012). The impact of pollution on worker productivity. *American Economic Review 102*(7), 3652–73.

Gruhl, H., N. Volkhausen, N. Pestel, and N. aus dem Moore (2025). Air pollution and the housing market: Evidence from Germany's Low Emission Zones. *Journal of Environmental Economics and Management 132*, 103161.

Gutiérrez, G. and S. Piton (2020). Revisiting the Global Decline of the (Non-housing) Labor Share. *American Economic Review: Insights 2*(3), 321–338.

Hanlon, W. W. (2020). Coal smoke, city growth, and the costs of the Industrial Revolution. *The Economic Journal 130*(626), 462–488.

Hanna, R. and P. Oliva (2015). The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City. *Journal of Public Economics 122*, 68–79.

Harari, M. (2020). Cities in bad shape: Urban geometry in India. *American Economic Review 110*(8), 2377–2421.

He, J., H. Liu, and A. Salvo (2019). Severe air pollution and labor productivity: Evidence from industrial towns in China. *American Economic Journal: Applied Economics 11*(1), 173–201.

Heissel, J. A., C. Persico, and D. Simon (2022). Does pollution drive achievement? The effect of traffic pollution on academic performance. *Journal of Human Resources 57*(3), 747–776.

Herrnstadt, E., A. Heyes, E. Muehlegger, and S. Saberian (2021). Air pollution and criminal activity: Microgeographic evidence from Chicago. *American Economic Journal: Applied Economics 13*(4), 70–100.

Hersbach, H., B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, and J.-N. Thépaut (2018). Era5 hourly data on single levels from 1959 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS).

Hoffmann, B. and J. P. Rud (2024). The unequal effects of pollution on labor supply. *Econometrica 92*(4), 1063–1096.

Holub, F. and B. Thies (2023). Air Quality, High-Skilled Worker Productivity and Adaptation: Evidence from GitHub. Available at SSRN.

INSEE (2023). French Economy Dashboard - 2023 Edition - Employment by business sector in 2023. [Available online at `https://www.insee.fr/en/outil-interactif/5543645/details/50_MTS/51_EPA/51D_figure4`].

Jégard, M. (2021). An Optimal Distribution of Polluting Activities Across Space. PSE Job Market Paper.

Kahn, M. E. and R. Walsh (2015). Cities and the Environment. *Handbook of Regional and Urban Economics 5*, 405–465.

Krebs, B. and S. Luechinger (2024). Air Pollution, Cognitive Performance, and the Role of Task Proficiency. *Journal of the Association of Environmental and Resource Economists*.

Künn, S., J. Palacios, and N. Pestel (2023). Indoor air quality and strategic decision making. *Management Science 69*(9), 5354–5377.

La Nauze, A. and E. Severnini (2025). Air pollution and adult cognition: Evidence from brain training. *Journal of the Association of Environmental and Resource Economists 12*(1), 221–255.

Le Boennec, R. and F. Salladarré (2017). The impact of air pollution and noise on the real estate market. The case of the 2013 European Green Capital: Nantes, France. *Ecological Economics 138*, 82–89.

Leroutier, M. and H. Ollivier (2025). The Cost of Air Pollution for Workers and Firms. CREST Working Paper No.05/March 2025.

Lichter, A., N. Pestel, and E. Sommer (2017). Productivity effects of air pollution: Evidence from professional soccer. *Labour Economics 48*, 54–66.

Liotta, C., V. Viguié, and Q. Lepetit (2022). Testing the monocentric standard urban model in a global sample of cities. *Regional Science and Urban Economics 97*, 103832.

Lyubich, E. (2025). The role of people versus places in individual carbon emissions. *American Economic Review 115*(5), 1439–1484.

Mills, E. S. (1967). An aggregative model of resource allocation in a metropolitan area. *American Economic Review 57*(2), 197–210.

Monte, F., S. J. Redding, and E. Rossi-Hansberg (2018). Commuting, migration, and local employment elasticities. *American Economic Review 108*(12), 3855–3890.

Moretti, E. (2013). Real wage inequality. *American Economic Journal: Applied Economics 5*(1), 65–103.

Qin, Y., J. Wu, and R. Zhang (2022). Can professional football players adapt to air pollution? Evidence from China. *Journal of Sports Economics 23*(3), 277–300.

Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy 90*(6), 1257–1278.

Rosen, S. (1979). Wage-based indexes of urban quality of life. *Current Issues in Urban Economics*, 74–104.

Rosenthal, S. S. and W. C. Strange (2008a). Agglomeration and hours worked. *Review of Economics and Statistics 90*(1), 105–118.

Rosenthal, S. S. and W. C. Strange (2008b). The attenuation of human capital spillovers. *Journal of Urban Economics 64*(2), 373–389.

van Donkelaar, A., M. S. Hammer, L. Bindle, M. Brauer, J. R. Brook, M. J. Garay, N. C. Hsu, O. V. Kalashnikova, R. A. Kahn, C. Lee, et al. (2021). Monthly global estimates of fine particulate matter and their uncertainty. *Environmental Science & Technology 55*(22), 15287–15300.

A Model Appendix

This Appendix provides extra information on the model presented in Section 2, including full expressions for equilibrium prices and quantities, as well as additional theoretical results relevant to the estimation of the overall effect of air pollution on these equilibrium prices and quantities.

A.1 Structural equations: fixed city surface

This section provides complementary expressions of equilibrium relationships in the case where city surface is held fixed. Expressions for equilibrium city population, wages and housing prices are given in the main text (Section 2.2). In complement, I show here the expression for equilibrium real wages, keeping city surface fixed. Expressions for the nominal wage premium and the real wage penalty are also given in the main text (Section 2.4). I also provide the expressions for the housing price premium when air pollution is controlled for (equation 32) and when it is not (equation 33), keeping spatial surface fixed.

Equilibrium real wages, fixed S I write real wages as equilibrium nominal wages, minus the product of the housing share in expenditure and housing price.

$$\begin{aligned}\ln(w^*_{\bar{S}}) - \beta \ln(q(0)^*_{\bar{S}}) = \Gamma_R + \Phi \left[\alpha\delta(1-\beta\kappa) \ln(\tilde{A}) - (1-\gamma-\alpha(1+\lambda)(1-\beta\kappa)) \ln(\theta) \right. \\ \left. - (\beta\alpha(1+\lambda)(1-\beta\kappa))(\ln(B) + \kappa \ln(T(\bar{x})) \right]\end{aligned} \tag{31}$$

As equation (31) shows, equilibrium real wages are higher where labour productivity $\tilde{A}$ and quality of life θ are higher; and all the more so if housing supply is sufficiently elastic (low κ) and if housing represents a low share of expenditure (low β). This is consistent with the Rosen-Roback intuition that real wages should negatively react to fundamental amenity levels. Under the very weak assumption that $\beta\kappa < 1$, real wages however decrease with construction productivity B.

Housing price response to city size, S fixed The housing price premium to higher city population estimated using an amenity shock writes:

$$\frac{\partial \ln(q(0)^*_{ct})}{\partial \ln(N^*_{ct})} = \frac{\eta^q_\theta}{\eta^N_\theta} = \kappa \, \frac{\alpha(1+\lambda)}{1-\gamma} \tag{32}$$

This housing price premium is again a positive function of agglomeration economies λ, as they attract a higher worker demand for housing. It is also a positive function of κ: the less elastic the housing supply, the higher the increase in housing price for a given shock.

(Mis)measurement of the housing price premium In the same fashion as in equations (18) and (20), the housing price premium to density is measured by regressing the housing price index on

city size or density. If one does not control for air pollution, one estimates the following:

$$\frac{d\ln(q(0)^*_{\bar{S},ct})}{d\ln(N^*_{\bar{S},ct})} = \frac{d\ln(q(0)^*_{ct})/d\ln(\theta_{ct})}{d\ln(N^*_{ct})/d\ln(\theta_{ct})} = \frac{\eta^q_\theta - \left[\eta^q_a z_a + \eta^q_\theta z_\theta\right]\phi_N \eta^N_\theta}{\eta^N_\theta\left[1-(\eta^q_a z_a + \eta^q_\theta z_\theta)\phi_N\right]} \tag{33}$$

Failing to control for air pollution implies that its deflating effect on housing prices (at the numerator) and its deflating effect on population (at the denominator) are not accounted for. This is extremely to the case of nominal wages in equation (18). Worse air quality from higher density attenuates (resp., inflates) the urban housing price premium if equilibrium housing prices react more (resp., less) to changes in local fundamentals than equilibrium city size does.

A.2 Structural equations: endogenous city surface

This subsection provides the full expressions for equilibrium city size, nominal wages, and housing prices, when the city surface S is allowed to vary, as functions of the model's parameters and fundamental city attributes: construction productivity $\tilde{B}$, labour productivity $\tilde{A}$ and amenity level θ. Shorter versions of these expressions are derived and discussed in Section 2.5 of the main text. $\tilde{\Phi}$ is equal to $(\tau - (1+\tau-2\beta\kappa)[(1-\gamma)(1+\delta)-\alpha(1+\lambda)\Phi])^{-1}$, and Γ_N, Γ_F and $\tilde{\Gamma}_H$ group constant terms. Appendix Table A.6 provides a calibration of the coefficients.

Equilibrium nominal wages, S free I follow the same procedure as in Section 2.2 with fixed city surface to obtain an expression for nominal wages with an endogenous city fringe.

$$\begin{aligned}\ln(w^*) &= \tilde{\Gamma}_F - \tilde{\Phi}\left[(1-\gamma-\alpha(1+\lambda))\left[2\beta\kappa\pi\ln(\theta) + (2\beta\kappa\pi + \beta\sigma\tau - \tau)\ln(\tilde{B})\right]\right. \\ &\quad \left. + \frac{\alpha}{1-\gamma}\left[\tau + (2\beta\kappa\pi - \tau)(1-\gamma)\delta\right]\ln(\tilde{A})\right] \\ &= \tilde{\Phi}\left[\tilde{\Gamma}_F - (1-\gamma-\alpha(1+\lambda))\left(\Gamma_\theta\ln(\theta) + \Gamma_B\ln(B)\right) + \frac{\alpha}{1-\gamma}\left(\tau + \delta(1-\gamma)\Gamma_A\right)\ln(\tilde{A})\right]\end{aligned} \tag{34}$$

Since $\Gamma_\theta/ > \Phi$ and $\alpha(1+\lambda) > 1-\gamma$, for the same improvement in consumption amenities θ, nominal wages increase more when the city fringe is allowed to react than when it does not. This is due to the fact as amenities improve, the expansion of the city allows for an even larger increase in population, and hence an even larger gain in labour productivity. Let us assume a positive shock on amenities as in Section 2.4. The inflow of inhabitants to the city triggers wage gains through agglomeration economies, but also raises local air pollution. Labour productivity and amenity levels decrease as a consequence. Equation (34) shows that the negative shock on θ will have an even stronger negative effect on wages if the city fringe is allowed to react. This effect is however mitigated by the spatial expansion of the city, which disperses pollutants and thus limits the increase in local air pollution.

In the main text, I denote:

$$\Gamma_\theta/\Phi = 2\beta\kappa - 1, \quad \Gamma_B/\Phi = \beta(2\beta\kappa + \tau(\sigma-1) - 1), \text{ and } \Gamma_A/\Phi = \beta\kappa(2-\tau)\frac{\alpha}{1-\gamma} - 1 \tag{35}$$

A.3 Theoretical implications for estimation: the effect of Z on w^* and $q(0)^*$

This subsection complements the results of Section 2.4 by discussing the model-derived expressions for the "exogenous" effect of local air pollution on equilibrium wages and housing prices, omitting its relation to city structure.

Air pollution and nominal wages Inserting equations (14), on the effects of Z_{ct}, and (15), on Z_{ct} formation, into equation (12) for equilibrium wages at the city level, the effect of local air pollution on nominal wages can be expressed as:

$$\frac{\partial w^*_{\bar{S},ct}}{\partial Z_{ct}} = ((1-\gamma-\alpha(1+\lambda))z_\theta - (\alpha(\delta+\beta\kappa))z_a)\Phi < 0 \tag{36}$$

Which, because $\alpha\lambda > 1-\alpha-\gamma$, as we show below, is unambiguously negative. In the presence of strong agglomeration economies, reduced-form estimates on the effect of air pollution on nominal wages cannot be used to infer whether any compensation for local air pollution does occur.

Air pollution and housing prices Doing the same as above, but looking at equilibrium housing prices from equation (13), I obtain:

$$\frac{\partial \ln(q(0)^*_{\bar{S},ct})}{\partial \ln(Z_{ct})} = -\alpha\kappa\left[(z_\theta + z_a) + (\lambda z_\theta + \delta z_a)\right]\Phi < 0 \tag{37}$$

Both agglomeration gains λ and congestion costs δ intervene to make the effect of Z_{ct} even more negative. As a consumption disamenity, higher air pollution drives workers away from the city, which depreciates housing prices, and even more so when the loss of workers implies a strong foregone productivity gain (a high λ). A similar mechanism arises due to local air pollution being a production disamenity ($-z_a < 0$): the loss of workers implies lower congestion, which depletes housing prices even more strongly if these costs matter a lot (a high δ).

Air pollution and real wages I measure real wages as equilibrium wages net of housing costs, and follow conventions by defining the latter as the product of the housing expenditure share β and the equilibrium housing price $q(0)^*_{\bar{S},ct}$. The effect of local air pollution on real wages thus writes:

$$\frac{\partial\left(\ln(w^*_{ct}) - \beta\ln(p^{H*}_{ct})\right)}{\partial\ln(Z_{ct})} = \Big[(1-\gamma-\alpha\,(1+\lambda)\,(1-\beta\kappa))z_\theta - \alpha\,\delta\,(1-\beta\kappa)\,z_a\Big]\Phi \tag{38}$$

The sign of the overall effect is here ambiguous. The productivity channel on the right-hand

side implies a negative effect ($-\alpha\ \delta(1-\beta\kappa)z_a < 0$ under the reasonable assumption that $\beta\kappa < 1$). However, the consumption amenity channel on the left-hand side induces a positive (resp., negative) impact of better amenities on real wages only if extra labour income from agglomeration economies, through $\alpha\lambda$, is larger (resp., smaller) than $\frac{1-\gamma}{1-\beta\kappa} - \alpha$. I discuss this strong condition further below when I study the measurement of the real wage differential of city size. In practice, housing supply is not sticky enough, i.e., κ is not large enough, for amenities to have a negative impact on real wages.

A.4 Model calibration

The discussion behind the choice of calibrated values of Table A.5 below is given in Section 6.1 of the main text.

Table A.5: Parameter values

Name	Notation	Value	Source
Labour share	α	0.75	Cette et al. (2019); Gutiérrez and Piton (2020)
Capital share	γ	0.2	
Housing share	β	0.33	Combes et al. (2021)
Inverse housing supply elasticity	κ	0.3	Combes et al. (2021); Baum-Snow and Duranton (2025)
Congestion	-δ	-0.4	Monte et al. (2018); Bilal (2023)
Agglomeration	λ	0.11	Deduced from TSLS estimation of (16)
Commuting disutility	τ	0.15	Combes et al. (2019)
Commuting cost within CBD	b	0.05	Combes et al. (2019)
Air pollution from population	ϕ_N	0.15	Estimation of (27), TSLS
Air pollution from footprint	-ϕ_S	-0.14	Estimation of (27), TSLS
Production disamenity	-ζ_a	-0.11	Deduced from estimation of (29), TSLS
Consumption disamenity	-ζ_θ	-0.06	Deduced from estimation of (29), TSLS

Notes: $\kappa = \frac{\sigma-1}{\sigma}$, where $\sigma > 1$ denotes the convexity of the cost of building taller, is an intensive-margin elasticity, hence the calibration to Combes et al. (2021). $\delta = 1/\epsilon = 1/2.5$ where ϵ is the dispersion parameter in Ahlfeldt et al. (2024); Monte et al. (2018); Bilal (2023).

Table A.6: Calibrated values of main coefficients of equilibrium city size, wages and housing prices

Equation	Coefficient	Symbolic value	Calibrated value
City size, S fixed	$\frac{d\ln N^*_{\bar{S}}}{d\ln\theta}$	$(1-\gamma)\Phi$	2.16
City size, S free	$\frac{d\ln N^*}{d\ln\theta}$	$2(1-\gamma)\beta\kappa\tilde{\Phi}$	2.29
City size, S fixed	$\frac{d\ln N^*_{\bar{S}}}{d\ln\tilde{A}}$	$\alpha(1-\beta\kappa)\Phi$	1.83
City size, S free	$\frac{d\ln N^*}{d\ln\tilde{A}}$	$(2-\tau)\alpha\beta\kappa\tilde{\Phi}$	1.98
Wages, S fixed	$\frac{d\ln w^*_{\bar{S}}}{d\ln\theta}$	$-(1-\gamma-\alpha(1+\lambda))\Phi$	0.09
Wages, S free	$\frac{d\ln w^*}{d\ln\theta}$	$-2(1-\gamma-\alpha(1+\lambda))\beta\kappa\tilde{\Phi}$	0.09
Wages, S fixed	$\frac{d\ln w^*_{\bar{S}}}{d\ln\tilde{A}}$	$\alpha(\delta+\beta\kappa)\Phi$	1.01
Wages, S free	$\frac{d\ln w^*}{d\ln\tilde{A}}$	$\frac{\alpha}{1-\gamma}\left[1-(2-\tau)(1-\gamma-\alpha(1+\lambda))\beta\kappa\tilde{\Phi}\right]$	1.02
Housing price, S fixed	$\frac{d\ln q(0)^*_{\bar{S}}}{d\ln\theta}$	$\kappa\alpha(1+\lambda)\Phi$	0.68
Housing price, S free	$\frac{d\ln q(0)^*}{d\ln\theta}$	$\frac{1}{\beta}\left(1-2Q\beta\kappa\tilde{\Phi}\right)$	0.54
Housing price, S fixed	$\frac{d\ln q(0)^*_{\bar{S}}}{d\ln\tilde{A}}$	$\kappa\alpha(1+\delta)\Phi$	0.85
Housing price, S free	$\frac{d\ln q(0)^*}{d\ln\tilde{A}}$	$\frac{\alpha}{\beta(1-\gamma)}\left(1-(2-\tau)Q\beta\kappa\tilde{\Phi}\right)$	0.68

Note: All calibrated values taken from Table A.5.

B Data Appendix

This Appendix contains additional information on the data sources and how I use them to construct wage and housing price outcomes.

B.1 Wage measure

For years 2002 to 2018, the period of study, the raw data contains about 57.8 million observations. Sample selection rather closely follows Combes et al. (2008), even though improvements in the quality of the data allow to impose less stringent constraints. For instance, I am able to keep part-time employees, since the number of hours worked is observed from 2002 on, as observed to the number of days worked.

Sample selection Table B.1 presents the evolution in the number of observations at some important steps of the sample selection. I start by dropping all observations with missing values in important variables. Given that the data is at the job level, some workers are observed several times in one year. In such a case, I select the observation with the highest number of hours worked. Then, I proceed to select private sector employees of working age (between 15 and 65 years old) who work in mainland France. This covers both full-time and part-time workers.[47] As I use soil characteristics as instruments for density, I exclude all workers of the agricultural sector in order to avoid endogeneity. Moreover, some sectors such as postal service or financial intermediation have specific reporting rules, which allow them to file their declaration at the regional level instead of the establishment level. This implies that for these industries, one cannot properly pinpoint the employment zone workers work in. As such, these industries are excluded from the sample. I also drop all (3-digit) industries that have less than 500 observations throughout the study period. In addition to public administration and agricultural activities, this implies excluding the following industries: financial services, insurance, postal services, telecommunications, coal extraction, uranium, thorium and other non-ferrous metal extraction, supporting activities in oil extraction and other extracting activities, manufacturing of fur products, coking, and manufacturing of magnetic and optical media. Finally, some outlying values in hourly wage, likely due to reporting mistakes in the number of hours worked or the yearly wage, are dropped by trimming the bottom 1% and top 1% of hourly wages for each year. The final sample consists of 25 million worker-year observations over the 2002-2018 period, out of which 21.6 million live and work in urban areas.

Accounting for unobserved sorting In order to filter out unobserved worker heterogeneity and localisation economies, I use the two-step method of Combes et al. (2008), and compute an "ad-

[47] Combes et al. (2008) use DADS data over the period 1976-1996, during most of which one could not observe the number of hours worked, but the number of *days* worked. As such, they focus on full-time employees.

[48] I thus drop Corsica and overseas *départements*.

[49] This includes agriculture, fishing and extractive industries.

[50] Regional reporting precludes the proper identification of the location of workers, which is crucial here. Industries that are allowed to use regional-level reporting include financial services, insurance, ICT, and postal services.

Table B.1: Number of observations for each step of the DADS sample selection, 2002-2018

Step	Description	Observations left
0	All observations	57,781,020
1	Mainland France[48] + wage and hours > 0	51,473,665
2	Only 1 observation per worker	51,463,451
3	Drop retirees, apprentices, elected officials	51,109,349
4	15-65 age range	51,091,059
5	Drop public sector and associations	37,153,371
6	Drop industries likely correlated with IV[49]	36,246,304
6	Drop industries with regional reporting[50]	28,063,799
7	Drop industries <500 observations	25,487,665
8	Drop top and bottom 1% outliers in hourly wage	24,977,890
9	Keep only those in selected urban areas	21,618,013

justed wage" W_{ct}^{adj}. I estimate this from the worker-year panel:

$$\ln w_{it} = W_{c(i),t}^{adj} + \overline{\text{PBLH}}_c + \lambda_s + \boldsymbol{X}_{\boldsymbol{c(i),s,t}}^{\boldsymbol{S}}\boldsymbol{\gamma}^{\boldsymbol{S}} + \boldsymbol{X}_{\boldsymbol{it}}^{\boldsymbol{W}}\boldsymbol{\gamma}^{\boldsymbol{W}} + \varepsilon_{it} \tag{39}$$

where w_{it} is the hourly wage of worker i for year t, $\overline{\text{PBLH}}_c$ is worker fixed effect and λ_s is a sector fixed effect. Exploiting the individual nature of the data allows to differentiate the role played by individual characteristics (age, age squared in X^W) from the role played by sector-city-specific characteristics X^S. X^S includes the share of executives[51] in local sector s in city c, the (log) share that each sector s represents in the total employment of city c and the (log) share that each sector s represents in the total number of establishments of city c, all at time t. W_{ct}^{adj} is a city-year fixed effect, and corresponds to the adjusted wage measure used as the main outcome of the wage analysis.

B.2 Housing price index

I construct a housing price index based on exhaustive transaction data from the DV3F (*Demande de Valeurs Foncières*) database, collected for tax collection purposes. The dataset is available from 2010 on, and covers the universe of housing and land transactions in France, excluding Alsace-Moselle, in the East. I focus on years 2010-2018 and select transactions of flats and houses, and thus exclude transactions of entire buildings or bare or agricultural land. As there are some outlying values due to errors in recording, I compute the price per sqm and trim the top and bottom 0.5% of prices per sqm at the *département*×year level. At this point, I have 5,186,887 transactions over 9 years. To ensure that the housing price index is computed on non-sparse cells, I select cities for which I observe at least 150 transactions per year, which brings the number of cities down from 403 to 257, and the number of transactions to 5,089,014.

As the empirical equivalent of the housing price gradient of the theoretical model, I compute

[51] I use the INSEE definition of occupations, which includes executives, engineers but also largely "independent" professional occupations such as lawyers, physicians and veterinarians (with employee status) or journalists.

the housing price index from the residual city-level variation χ_c in a regression that I run at the transaction level. For each year t, I regress the log price of transaction j of a dwelling located in city c and sold during quarter q:

$$\begin{aligned}\ln(\text{transaction price}_{j(c)}) = \alpha_0 + \alpha_1 \ln(\text{built area}_j) + \alpha_2 \ln(1 + \text{plot area}) + \alpha_3 \text{ distance CBD}_j \\ + \alpha_4 \mathbb{1}\{\text{house}_j\} + \alpha_5 \text{ building period}_j + \omega_q + \chi_c + u_j\end{aligned} \tag{40}$$

where I regress the log of the amount of the transaction on the log built area, the log plot area (+1), distance in km to the CBD, a dummy equal to one if the good sold is a house (0 if it's a flat), a factor of 6 construction periods, and a quarter fixed effect. χ_c thus denotes the city fixed effect that measures the housing price index at the CBD, net of differences in housing composition, and is my measure of $\ln(\widehat{q(0)})_{ct}$.

For the sake of exposition, Table B.2 provides estimations of several variants of equation (40), grouping all years together and replacing χ_c by a city-year fixed effect (not shown). In particular, comparing Column (2) to Column (3), (5) or (6) shows that using level distance to the CBD, log distance to the CBD, or a quadratic form as a control for the within-city housing price gradient has no incidence. It also shows that the estimations are not sensitive to the exclusion of transactions that occur in the CBD (and for which distance is equal to 0).

Table B.2: Hedonic regressions of housing prices and distance to CBD

	Outcome: log transaction value					
	(1)	(2)	(3)	(4)	(5)	(6)
Distance to CBD (km)		-0.0180***		-0.0190***	-0.0186**	
		(0.0029)		(0.0023)	(0.0078)	
Distance to CBD (km) squared					1.28×10^{-5}	
					(9.31×10^{-5})	
log (distance to CBD + 0.01)			-0.0218***			
			(0.0053)			
log distance to CBD						-0.2935***
						(0.0850)
log built area	0.7590***	0.7581***	0.7565***	0.7492***	0.7582***	0.7550***
	(0.0080)	(0.0099)	(0.0083)	(0.0183)	(0.0100)	(0.0218)
log (plot area + 1)	0.0130*	0.0346***	0.0195***	0.0349***	0.0348***	0.0323***
	(0.0070)	(0.0026)	(0.0064)	(0.0032)	(0.0032)	(0.0043)
House dummy	-0.0246	0.0500*	-0.0015	0.0064	0.0511***	0.0210
	(0.0488)	(0.0262)	(0.0494)	(0.0275)	(0.0190)	(0.0146)
Q2	0.0080***	0.0083***	0.0077***	0.0084***	0.0083***	0.0083***
	(0.0009)	(0.0011)	(0.0009)	(0.0016)	(0.0010)	(0.0018)
Q3	0.0364***	0.0350***	0.0356***	0.0382***	0.0350***	0.0385***
	(0.0029)	(0.0027)	(0.0028)	(0.0033)	(0.0026)	(0.0036)
Q4	0.0157***	0.0176***	0.0157***	0.0190***	0.0176***	0.0185***
	(0.0012)	(0.0015)	(0.0012)	(0.0016)	(0.0014)	(0.0016)
Built before 1914	0.0512	0.0267	0.0417	0.0360	0.0258	-0.0077
	(0.0870)	(0.0655)	(0.0839)	(0.0854)	(0.0583)	(0.0469)
Built 1914-1944	0.1315**	0.0345	0.1068*	0.0437	0.0335**	0.0363**
	(0.0610)	(0.0214)	(0.0589)	(0.0308)	(0.0148)	(0.0173)
Built 1945-1960	0.0743**	0.0053	0.0499*	0.0225	0.0042	0.0099
	(0.0302)	(0.0132)	(0.0285)	(0.0195)	(0.0075)	(0.0114)
Built 1961-1974	0.0184**	-0.0010	0.0130	0.0208**	-0.0014	0.0113
	(0.0086)	(0.0070)	(0.0084)	(0.0097)	(0.0077)	(0.0141)
Built 1975-1989	0.1110***	0.1110***	0.1111***	0.1201***	0.1108***	0.1055***
	(0.0135)	(0.0136)	(0.0147)	(0.0162)	(0.0146)	(0.0262)
Built 1990-2012	0.1883***	0.2067***	0.1962***	0.2019***	0.2069***	0.1902***
	(0.0182)	(0.0151)	(0.0183)	(0.0193)	(0.0143)	(0.0248)
Built after 2013	0.2019***	0.2316***	0.2179***	0.2000***	0.2320***	0.1915***
	(0.0231)	(0.0175)	(0.0231)	(0.0214)	(0.0161)	(0.0255)
Including distance CBD = 0	X	X	X		X	
City×year fixed effects	X	X	X	X	X	X
Observations	5,089,014	5,089,014	5,089,014	3,651,168	5,089,014	3,651,168
R^2	0.57379	0.63560	0.57919	0.59338	0.63562	0.58951
Within R^2	0.46621	0.54361	0.47296	0.51340	0.54365	0.50877

Sources: DV3F, IGN.

Notes: Standard errors clustered at the MSA level in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

The top panel displays variations of the control for distance to the CBD, while the bottom panel displays controls common to all regressions (built area, plot area, house dummy, quarter of transaction and construction date). Columns 1-3 and 5 include transactions located in the CBD (defined as the central municipality of the city), while Columns 4 and 6 don't. All columns include city×year fixed effects, which constitute the measure of the housing price index used in the main analysis.

B.3 Summary statistics

Table B.3: Summary statistics - French MSAs, 2002-2018

Variable	Mean	Std Dev.	Min.	Max.	N×T
Main variables					
Hourly log wage W^{raw} (2019 €)	2.28	0.10	1.97	2.61	6,069
Adjusted log wage W^{adj}	0.30	0.26	-0.24	0.96	6,069
Employment density	169.9	124.32	25.53	1,121	6,069
Log employment density	4.93	0.64	3.24	7.02	6,069
Population density	189.0	140.6	27.55	1,264	6,069
Log population density	5.03	0.65	3.32	7.14	6,069
Population (thousands)	140.9	674.2	2.67	12,481	6,069
Log population	10.73	1.21	7.89	16.34	6,069
Area (sq. km)	618.1	1,190	15.52	17,284	6,069
Log area	5.71	1.17	2.74	9.76	6,069
$PM_{2.5}$ concentration ($\mu g/m^3$)	12.52	2.19	6.90	33.2	6,069
Log PM concentration	2.51	0.17	1.93	3.50	6,069
Number of $PM_{2.5}$ pixels	729	1,425	18	21,136	6,069
Housing price index (€/km^2)	-0.33	0.33	-1.15	1.18	1,700
Weather controls					
Mean winter temperature (°C)	7.03	2.17	-7.35	12.27	6,069
Mean summer temperature (°C)	17.14	1.84	7.25	23.15	6,069
Mean wind speed (km/h)	13.52	3.98	1.18	26.81	6,069
Mean cloud cover (%)	61.81	6.42	3.78	73.50	6,069
Total precipitation (m)	0.83	0.23	0.30	2.01	6,069
Geographical controls					
Mean elevation (m)	248.45	293.27	0.94	2,286	6,069
Ruggedness (m)	454.44	581.20	2.2	4,030	6,069
Sea (dummy, %)	17.93				6,069
Waterway (dummy, %)	38.41				6,069
Mountainous (dummy, %)	31.42				6,069
Instruments					
Log 1806 population	9.78	1.35	0	13.83	6,069
Topsoil CEC: medium (%)	23.2				6,069
Some soil differentiation (%)	40.9				6,069
Share of swelling minerals	19.2	0.25	0	1	6,069
Share of clay/silts	33.8	0.30	0	1	6,069
Number of days with low PBLH	104.1	15.6	6	212	6,069

Sources: INSEE, DADS, ACAG, ERA5, IGN, Observatoire des Territoires, Cassini-EHESS, DV3F and ESDAC.

C Instrumental variables

C.1 Instrumenting for cross-city variation in density

Requirements from the model As shown in Section 2.4, an instrument for city size must shift consumption amenities θ rather than labour productivity $\tilde{A}$, which affects wages directly. When housing prices, and hence real wages, are the outcome, it must not shift construction productivity $\tilde{B}$ either, as the latter affects housing prices directly. In addition, the decomposition of the wage–density gradient in Section 4 requires the instruments for city structure to affect $PM_{2.5}$ only through density and surface area. This section describes each soil instrument, the channel through which it is expected to affect density, and the channels through which it could violate these requirements.

Construction of the soil instruments All three instruments are drawn from the European Soil Database (ESDB), whose attributes are expressed as categorical classes on a 1 km $\times$ 1 km grid. They are derived from the soil typological units of the 1:1,000,000 European soil map and contain no contemporary measurement. For each city, I compute the share of its surface area that falls into a given class. The instruments enter the first stages with quadratic terms, as I expect the relationship between soil quality and current city size is expected to be hump-shaped. Figure C.2 shows this hump shape for each instrument, and the first stages in Tables C.1 and C.2 confirm it conditional on controls.

Subsoil cation exchange capacity Cation exchange capacity (CEC) measures the soil's ability to retain positively charged nutrients such as calcium (Ca^{2+}), magnesium (Mg^{2+}) and potassium (K^{+}), and is therefore a central determinant of chemical fertility. The instrument is the share of the city's surface area with a medium subsoil CEC. Medium values correspond to soils that retain enough nutrients to sustain cultivation without the heavy, hard-to-work textures associated with the highest values, which were costly to till before mechanisation. The intended channel is historical: such soils supported agricultural surpluses, and hence early settlement, which persists in current city size. The main threat to validity is a construction-cost channel. CEC rises with clay content and with the share of swelling clay minerals, and clay-rich subsoils are those most exposed to shrink–swell movements, which damage foundations and raise construction costs. This would make subsoil CEC partly a shifter of $\tilde{B}$, which remains a valid instrument for nominal wages, but not for housing prices or real wages. The restriction to the medium class limits this concern, as shrink–swell hazard concentrates in the most clay-rich soils.

Medium depth to a gleyed horizon A gleyed horizon is a soil layer saturated with water for long periods, in which water replaces oxygen. When it lies close to the surface, roots lack oxygen and land is hard to work. When it lies very deep, soils tend to dry out. Land with a gleyed horizon at a medium depth is therefore seasonally moist but workable, typically in valley bottoms and

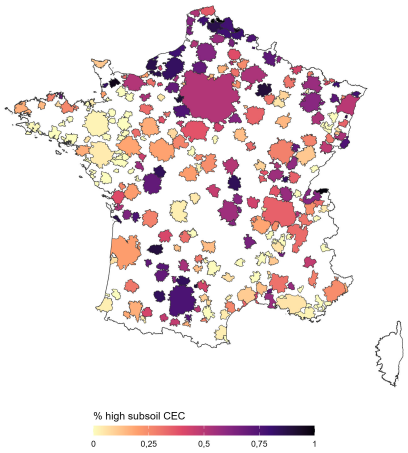
alluvial plains, and was well suited to pre-industrial agriculture, both for crops and for meadows. The instrument is the share of the city's surface area with a gleyed horizon at a medium depth. Two historical channels make it relevant for settlement: agricultural workability before modern drainage, and the disease environment, as permanently waterlogged land was associated with endemic malaria in several French regions until the nineteenth century. Both operate through historical amenities and agricultural potential, rather than through current productivity. The main threat is again a construction-cost channel: high water tables raise foundation, drainage and basement costs, and correlate with flood exposure. As for CEC, this would invalidate the instrument for housing prices and real wages, but not for nominal wages. The restriction to the medium class again limits the concern, as construction costs are highest where the gleyed horizon is shallowest.

High topsoil organic carbon The third instrument is the share of the city's surface area whose topsoil organic carbon content falls in the high or very high classes. Organic carbon improves water retention and nutrient cycling, but high contents also reflect slow decomposition, typically under grassland and in cool or moist conditions. The instrument thus captures land better suited to grazing and meadows than to intensive arable farming, which shaped the historical density of settlement. Unlike the two other instruments, organic carbon is derived in the ESDB through expert pedotransfer rules that combine soil type, climate and the typical land use associated with each soil typological unit in the 1:1,000,000 soil map (**?**). It therefore contains no local measurement and does not reflect actual land cover around cities. The construction-cost channel is thus very limited in its case.

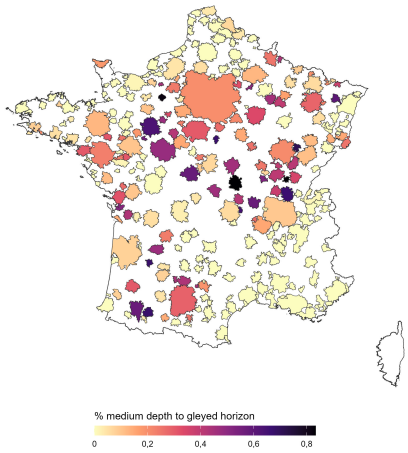
Productivity channel For all three instruments, a direct effect on current labour productivity would run through agriculture, which now accounts for about 2% of employment and value added. Agricultural workers are excluded from the wage sample, and the two-step adjusted wage nets out sectoral composition. A remaining channel could run through agri-food industries, if soil quality raises their local productivity, for instance through cheaper or better inputs.

Figure C.1: Geological instruments and urban density

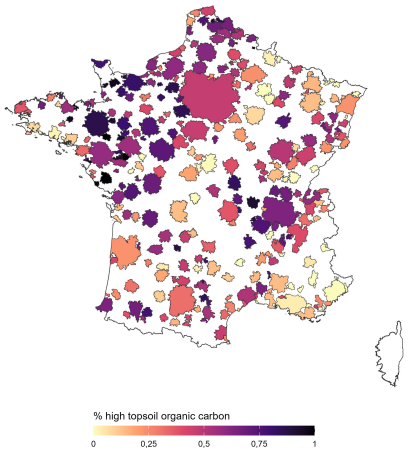
(a) High subsoil CEC



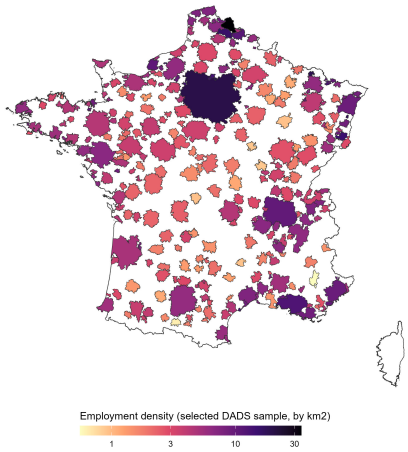
(b) Medium depth to gleysoil



(c) High topsoil OCC



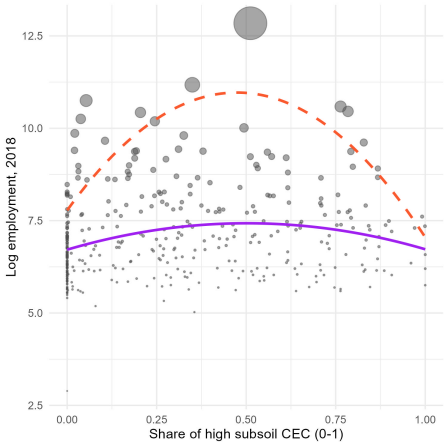
(d) Employment density



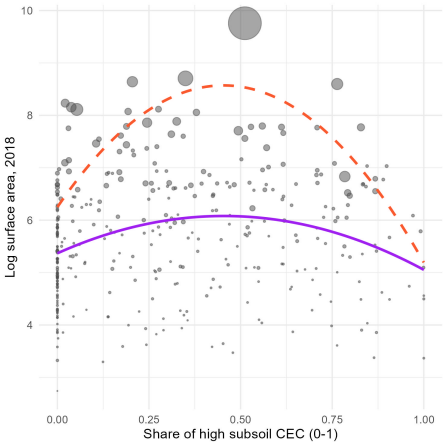
Sources: INSEE, ESDAC, DADS.
Notes: All values are averages over the 2006-2018 sample period. Values of Panels (a), (b) and (c) are computed as shares of the total surface area of the city. Panel (d) measures employment density as the number of workers in the DADS final sample, divided by the total surface area of the city.
Paris and south-western cities, such as Toulouse, compound both high CEC, a medium depth to a gleyed soil horizon and relatively low topsoil OCC, but have different levels of density. Cities located on the Mediterranean seafront, like Marseille or Nice in the south east, are very dense but do not have any of these soil fertility advantages. Central cities have high shares of medium depth to gleysoil but comparatively lower density. Northern cities display high subsoil CEC and topsoil OCC, but very deep gleyed soil (hence the low medium depth). The instruments are thus not perfectly correlated.

Figure C.2: City size and surface area peak at intermediate soil fertility

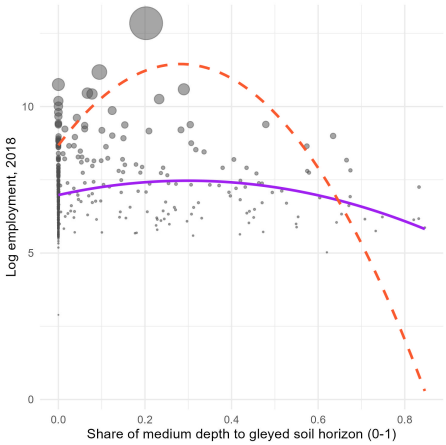
(a) Subsoil CEC and employment



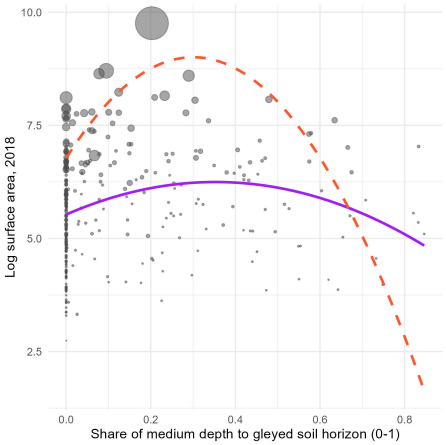
(b) Topsoil CEC and area



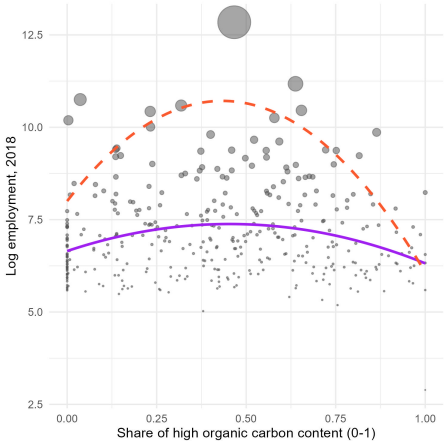
(c) Depth to gleysoil and employment



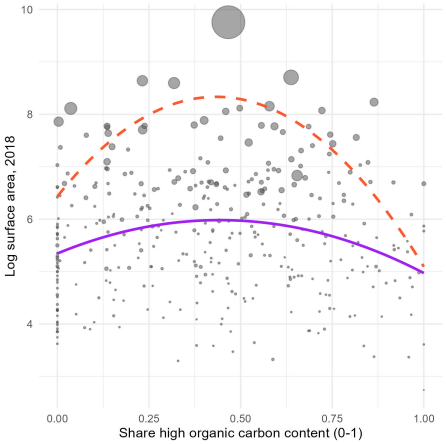
(d) Depth to gleysoil and area



(e) Topsoil OCC and employment



(f) Topsoil OCC and area



Sources: ESDB, DADS, IGN.
Notes: Each row shows a different soil fertility feature on the x-axis. The left column shows employment on the y-axis, and the right column shows surface area on the y-axis. Each point corresponds to a city, and the size of the point is proportional to city employment. The plain purple curve fits a quadratic function to the points, without weighting by city employment, while the red dashed line fits a quadratic function, weighting by city employment.

C.2 Instrumenting for cross-city variation in air quality

The planetary boundary layer (PBL) height instrument counts, for each city, the days on which daily PBL height is an anomaly, i.e., when it falls below a benchmark. The choice of benchmark is not innocuous. Let PBLH_{cd} denote daily PBL height in city c on day d, $\overline{\text{PBLH}}_c = \mathbb{E}[\text{PBLH}_{cd}]$ its long-term mean, and $\overline{\text{PBLH}}$ the common (pooled) mean across cities. Comparing to the city's own mean counts days with $\text{LowPBLH}_{cd} \equiv \text{PBLH}_{cd}/\overline{\text{PBLH}}_c - 1 < -(1-a)$, while comparing to the common mean counts days with $\widetilde{\text{PBLH}}_{cd} \equiv \text{PBLH}_{cd}/\overline{\text{PBLH}} - 1 < -(1-a)$, and these are not interchangeable.

A common benchmark applies, in own-mean terms, a *city-specific* cut-off:

$$\widetilde{\text{PBLH}}_{cd} < -(1-a) \iff \text{PBLH}_{cd} < a\,\overline{\text{PBLH}} \iff \text{LowPBLH}_{cd} < a\,\frac{\overline{\text{PBLH}}}{\overline{\text{PBLH}}_c} - 1 \equiv \tau_c.$$

The cut-off $\tau_c = a\,(\overline{\text{PBLH}}/\overline{\text{PBLH}}_c) - 1$ is a deterministic, decreasing function of the city's level $\overline{\text{PBLH}}_c$. A city with a chronically shallow PBL ($\overline{\text{PBLH}}_c < \overline{\text{PBLH}}$) faces a lax cut-off and is counted as low-PBL most of its days. A city with a chronically thick one ($\overline{\text{PBLH}}_c > \overline{\text{PBLH}}$) faces a stringent one and is almost never counted. With $a = 0.7$, a city at $\overline{\text{PBLH}}_c = 0.5\,\overline{\text{PBLH}}$ has $\tau_c = +0.40$ and one at $\overline{\text{PBLH}}_c = 1.5\,\overline{\text{PBLH}}$ has $\tau_c = -0.53$. The common-benchmark count is therefore essentially a monotone transform of $\overline{\text{PBLH}}_c$, which recovers the topographic level, i.e., elevation and coastal position, which also loads on local amenities. Own-mean normalisation removes the level by construction, leaving cross-city variation in the *shape* of each city's centred distribution, by measuring how often PBL height dips below its own norm.

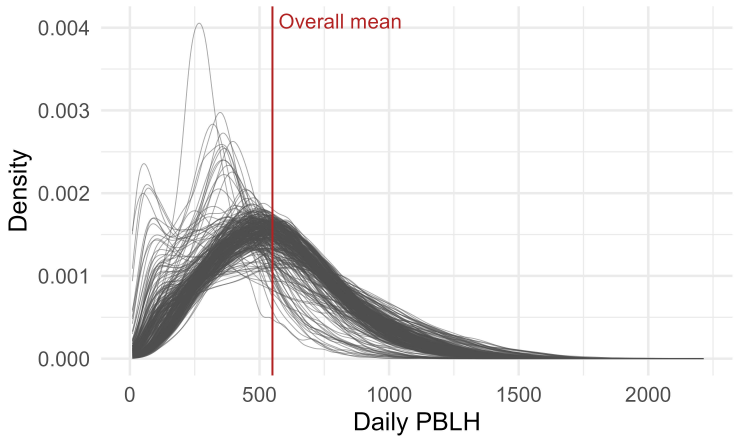
Figure C.3 makes this visible: the cross-city dispersion in $\overline{\text{PBLH}}_c$ is large (Panel a), so centring on the common mean (Panel b) merely shifts each city's whole distribution by its own level and the count IV still tracks that level, whereas own-mean normalisation (Panel c) centres every city at zero and leaves only tail-shape differences.

The factor a trades off two forces. A low a (a large required drop) selects only the days of most extreme vertical stagnation, where the instrument's relevance is cleanest; but set too low, some cities never reach the cut-off and take a zero value, so identification comes from a subset of cities and estimates a very local effect. A high a (a small required drop) mechanically counts many days in every city, so the counts saturate and cross-city variation (and with it instrument strength) wavers. The value $a = 0.7$ is an empirically sound middle ground. Table **??** reports first stages across a range of a.

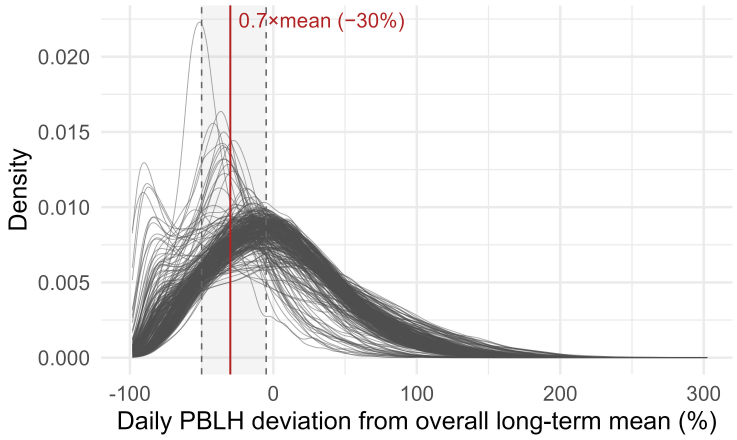
Two balance checks support treating the own-mean count as orthogonal to fixed geography. First, regressing $\overline{\text{PBLH}}_c$ on coarse geography (elevation, coastal distance, terrain, region) returns a high R^2: the level is geography. Second, regressing the own-mean count on the same geography (and on the recovered amenity θ_i) returns a weak association: the surviving shape variation is not. Table **??** reports both.

Figure C.3: Daily PBL height, overall means, and city-specific means

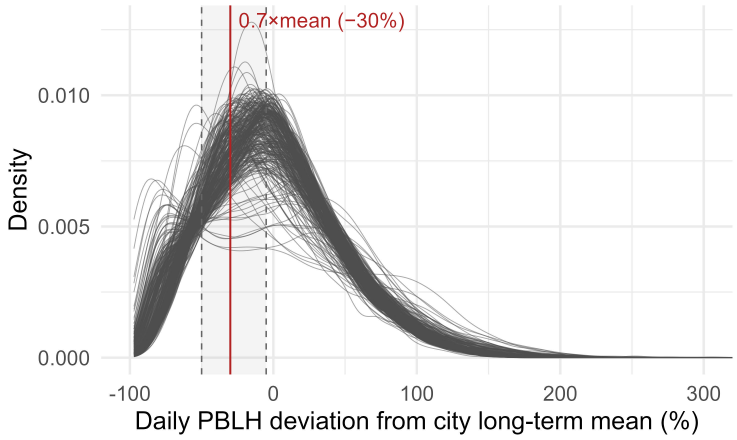
(a) No centring (levels)



(b) Centring around overall mean ($\overline{\text{PBLH}} = 550$)



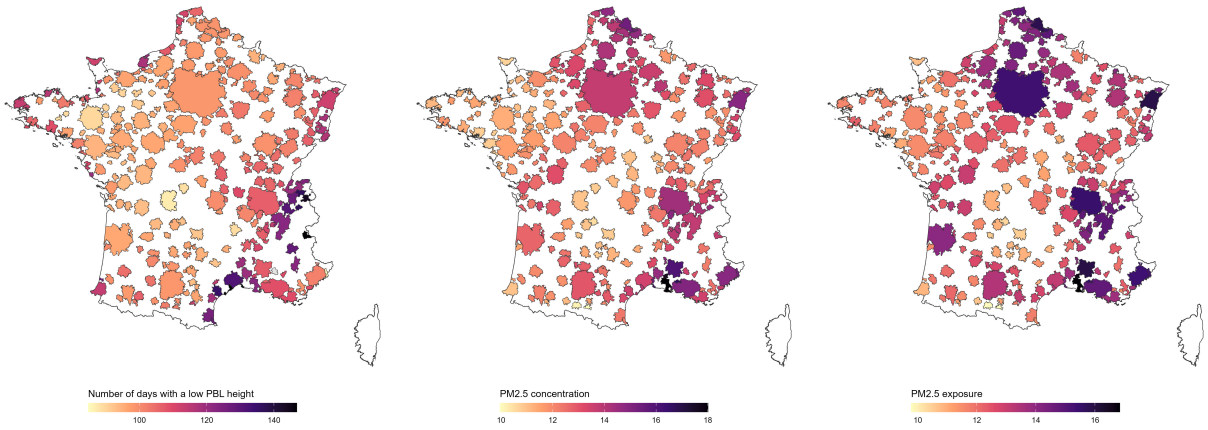
(c) Centring around city-specific mean ($\overline{\text{PBLH}}_c$)



Sources: ERA5, INSEE, IGN.
Notes: Each line is one city's density of daily PBL height (2002–2018); hourly ERA5 measures are averaged over the city surface and aggregated to the daily level. Panel (a) plots levels, with the vertical line at the pooled mean across all cities and days ($\overline{\text{PBLH}} = 550$ m). Panel (b) expresses each day as a deviation from that common mean, $\text{PBLH}_{cd}/\overline{\text{PBLH}} - 1$; Panel (c) as a deviation from the city's own mean, $\text{PBLH}_{cd}/\overline{\text{PBLH}}_c - 1$. In (b) and (c) the solid line marks the $a = 0.7$ cut-off (-30%) and the shaded area spans deviations of 5% to 50% below the benchmark ($a = 0.95$ to $a = 0.5$).

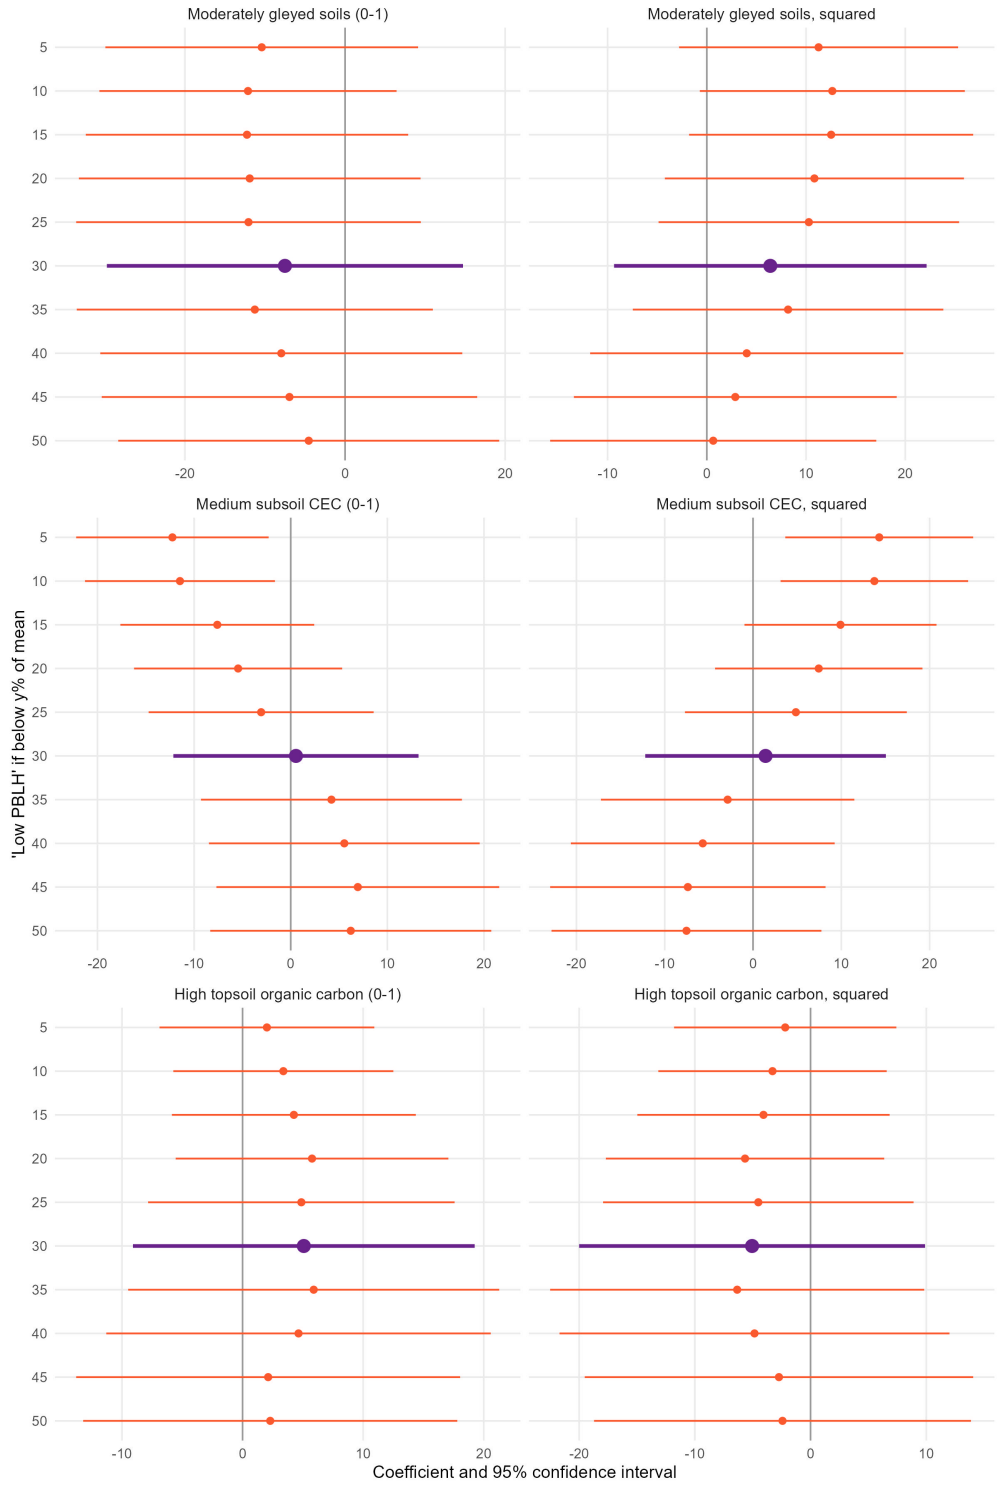
Figure C.4: Raw PBL height instrument, $PM_{2.5}$ concentration and $PM_{2.5}$ exposure

(a) Days with a low PBL height (b) $PM_{2.5}$ concentration (c) $PM_{2.5}$ exposure



Sources: ERA5, ACAG, INSEE, IGN.
Notes: All values are averages over the 2006-2018 sample period. $PM_{2.5}$ is measured by averaging pixels over the city's surface area. $PM_{2.5}$ exposure is measured by taking a population-weighted average of $PM_{2.5}$ concentration at the neighbourhood level, and is the main measure of air pollution used throughout the paper.

Figure C.5: No correlation between instruments for city structure and for air pollution



Sources: ESDB, ERA5, IGN, INSEE.
Note: Point estimates of regressions of the PBL height IV (for local air pollution) on the soil IV (for density) and corresponding 95% confidence intervals, at different levels of factor a in equation (3.4). The main regression results use $a = 0.7$, shown in bold and purple.

C.3 First-stage regression results

Table C.1 shows first-stage regression results for city size. Both soil instruments and 1806 population predict current population, while 1836 density in the central municipality of the city and the PBLH IV (intended to predict air pollution) don't. This holds controlling for weather (Columns 5-6) or not (Columns 1-4), and taking each instrument independently (Columns 1-3) or together (Columns 4-6).

Table C.2 does the same for city surface, and shows the relevance of both lagged population and soil instruments, while the instrument for air pollution does not predict it. Importantly, the expected inverted U-shaped relationship between soil fertility and city size and surface is confirmed (Columns 2-6 of Tables C.1 and C.2).

Table C.3 provides first-stage regression results for air pollution.

Table C.1: First-stage regression results - City size

	Outcome variable: log employment, excl. missing $\mathbf{q}(\mathbf{0})$					
	(1)	(2)	(3)	(4)	(5)	(6)
log 1806 population	0.641***			0.624***	0.600***	0.603***
	(0.090)			(0.085)	(0.089)	(0.090)
log 1836 central density	0.102			0.082	0.105	0.104
	(0.063)			(0.061)	(0.064)	(0.064)
% medium gleysoil depth		3.47***		2.01***	1.62**	1.60**
		(0.970)		(0.689)	(0.680)	(0.680)
Sq. % medium gleysoil depth		-5.17***		-3.11***	-2.71***	-2.70***
		(1.55)		(0.990)	(0.984)	(0.980)
% high organic carbon			1.65**	1.11*	1.36**	1.28**
			(0.692)	(0.593)	(0.597)	(0.609)
Sq. % high organic carbon			-1.22	-0.837	-1.11	-1.03
			(0.753)	(0.703)	(0.720)	(0.734)
Nb days low PBL height						-0.001
						(0.003)
Observations	2,223	2,241	2,241	2,223	2,223	2,214
Number of clusters (cities)	247	249	249	247	247	246
R^2	0.57	0.24	0.23	0.60	0.61	0.61
F-stat	957.22	58.09	36.92	357.35	354.49	349.33
	Outcome variable: log employment, incl. missing $\mathbf{q}(\mathbf{0})$					
	(1)	(2)	(3)	(4)	(5)	(6)
log 1806 population	0.594***			0.581***	0.582***	0.585***
	(0.076)			(0.072)	(0.075)	(0.075)
log 1836 central density	0.081			0.069	0.065	0.064
	(0.057)			(0.056)	(0.058)	(0.058)
% medium gleysoil depth		3.24***		2.31***	1.99***	1.97***
		(0.904)		(0.650)	(0.634)	(0.635)
Sq. % medium gleysoil depth		-4.99***		-3.70***	-3.33***	-3.32***
		(1.51)		(0.978)	(0.952)	(0.951)
% high organic carbon			1.78***	0.927*	1.12**	1.07**
			(0.612)	(0.516)	(0.514)	(0.525)
Sq. % high organic carbon			-1.71***	-0.894	-1.09*	-1.04*
			(0.651)	(0.567)	(0.575)	(0.586)
Nb days low PBL height						-0.002
						(0.002)
Observations	5,066	5,134	5,134	5,066	5,066	5,049
Number of clusters (cities)	298	302	302	298	298	297
R^2	0.565	0.275	0.260	0.586	0.602	0.601
F-stat	1,884	114.48	60.58	704.45	710.48	700.50
	Outcome variable: log population, incl. missing $\mathbf{q}(\mathbf{0})$					
	(1)	(2)	(3)	(4)	(5)	(6)
log 1806 population	0.567***			0.553***	0.534***	0.536***
	(0.077)			(0.073)	(0.075)	(0.075)
log 1836 central density	0.104*			0.093	0.111*	0.110*
	(0.059)			(0.057)	(0.059)	(0.059)
% medium gleysoil depth		3.34***		2.40***	2.11***	2.10***
		(0.901)		(0.652)	(0.641)	(0.642)
Sq. % medium gleysoil depth		-5.20***		-3.89***	-3.60***	-3.59***
		(1.53)		(1.00)	(0.980)	(0.980)
% high organic carbon			1.96***	1.04**	1.22**	1.18**
			(0.605)	(0.505)	(0.507)	(0.517)
Sq. % high organic carbon			-1.93***	-1.06*	-1.21**	-1.17**
			(0.648)	(0.568)	(0.578)	(0.588)
Nb days low PBL height						-0.003
						(0.002)
Observations	4,172	4,228	4,228	4,172	4,172	4,158
Number of clusters (cities)	298	302	302	298	298	297
R^2	0.567	0.284	0.270	0.592	0.602	0.602
F-stat	1,524	102.93	59.98	580.29	578.25	569.23
Geography controls	X	X	X	X	X	X
Weather controls					X	X
Mean PBLH control						X
Year fixed effects	X	X	X	X	X	X

Notes: Geography controls include mean elevation, ruggedness, share of land covered by mountains, and dummies for sea access and waterway presence. Weather controls include summer temperature, winter temperature, mean wind speed and total precipitation.

Table C.2: First-stage regression results - City surface area

	Outcome variable: log surface area, excl. missing $\mathbf{q(0)}$					
	(1)	(2)	(3)	(4)	(5)	(6)
log 1806 population	0.713***			0.685***	0.667***	0.646***
	(0.077)			(0.071)	(0.077)	(0.078)
log 1836 central density	0.102			0.082	0.105	0.104
	(0.063)			(0.061)	(0.064)	(0.064)
% medium gleysoil depth		3.47***		2.01***	1.62**	1.60**
		(0.970)		(0.689)	(0.680)	(0.680)
Sq. % medium gleysoil depth		-5.17***		-3.11***	-2.71***	-2.70***
		(1.55)		(0.990)	(0.984)	(0.980)
% high organic carbon			1.63**	1.17*	1.31**	1.36**
			(0.755)	(0.654)	(0.664)	(0.678)
Sq. % high organic carbon			-1.50*	-1.17	-1.34*	-1.41*
			(0.827)	(0.760)	(0.776)	(0.792)
Nb days low PBL height						-0.001
						(0.003)
Observations	2,223	2,241	2,241	2,223	2,223	2,214
Number of clusters (cities)	247	249	249	247	247	246
R^2	0.57	0.24	0.23	0.60	0.61	0.61
F-stat	957.22	58.09	36.92	357.35	354.49	349.33
	Outcome variable: log surface area, incl. missing $\mathbf{q(0)}$					
	(1)	(2)	(3)	(4)	(5)	(6)
log 1806 population	0.594***			0.581***	0.582***	0.585***
	(0.076)			(0.072)	(0.075)	(0.075)
log 1836 central density	0.081			0.069	0.065	0.064
	(0.057)			(0.056)	(0.058)	(0.058)
% medium gleysoil depth		3.24***		2.31***	1.99***	1.97***
		(0.904)		(0.650)	(0.634)	(0.635)
Sq. % medium gleysoil depth		-4.99***		-3.70***	-3.33***	-3.32***
		(1.51)		(0.978)	(0.952)	(0.951)
% high organic carbon			1.78***	0.927*	1.12**	1.07**
			(0.612)	(0.516)	(0.514)	(0.525)
Sq. % high organic carbon			-1.71***	-0.894	-1.09*	-1.04*
			(0.651)	(0.567)	(0.575)	(0.586)
Nb days low PBL height						-0.002
						(0.002)
Observations	5,066	5,134	5,134	5,066	5,066	5,049
Number of clusters (cities)	298	302	302	298	298	297
R^2	0.565	0.275	0.260	0.586	0.602	0.601
F-stat	1,884	114.48	60.58	704.45	710.48	700.50
Geography controls	X	X	X	X	X	X
Weather controls					X	X
Mean PBLH control						X
Year fixed effects	X	X	X	X	X	X

Notes: Geography controls include mean elevation, ruggedness, share of land covered by mountains, and dummies for sea access and waterway presence. Weather controls include summer temperature, winter temperature, mean wind speed and total precipitation.

Table C.3: First-stage regression results - City $PM_{2.5}$ exposure

	Outcome variable: log $PM_{2.5}$ "exposure", excl. missing $\mathbf{q}(\mathbf{0})$					
	(1)	(2)	(3)	(4)	(5)	(6)
% days low PBL height	0.578***	0.392***	0.405***	0.332***	0.391***	0.353***
	(0.148)	(0.120)	(0.120)	(0.115)	(0.119)	(0.114)
log 1806 population			0.009			0.010
			(0.009)			(0.009)
log 1836 central density			0.006			0.009
			(0.007)			(0.007)
% medium gleysoil depth				-0.242***		-0.268***
				(0.064)		(0.066)
Sq. % medium gleysoil depth				0.231**		0.265***
				(0.100)		(0.102)
% high organic carbon					-0.031	-0.033
					(0.058)	(0.054)
Sq. % high organic carbon					0.041	0.041
					(0.064)	(0.061)
Observations	2,223	2,241	2,241	2,223	2,223	2,214
Number of clusters (cities)	247	249	249	247	247	246
R^2	0.749	0.802	0.808	0.814	0.802	0.822
F-stat	105.41	55.68	34.23	66.66	19.56	39.73
	Outcome variable: log $PM_{2.5}$ "exposure", incl. missing $\mathbf{q}(\mathbf{0})$					
	(1)	(2)	(3)	(4)	(5)	(6)
% days low PBL height	0.634***	0.435***	0.445***	0.383***	0.434***	0.399***
	(0.132)	(0.109)	(0.108)	(0.105)	(0.108)	(0.103)
log 1806 population			0.010			0.010
			(0.007)			(0.007)
log 1836 central density			0.005			0.007
			(0.007)			(0.006)
% medium gleysoil depth				-0.237***		-0.252***
				(0.058)		(0.058)
Sq. % medium gleysoil depth				0.229**		0.248***
				(0.090)		(0.090)
% high organic carbon					-0.017	-0.033
					(0.048)	(0.045)
Sq. % high organic carbon					0.021	0.040
					(0.053)	(0.049)
Observations	3,900	3,900	3,848	3,900	3,900	3,848
Number of clusters (cities)	300	300	296	300	300	296
R^2	0.724	0.775	0.781	0.788	0.775	0.795
F-stat	210.87	110.24	66.02	116.44	37.26	67.96
	Outcome variable: log raw $PM_{2.5}$ mean, excl. missing $\mathbf{q}(\mathbf{0})$					
	(1)	(2)	(3)	(4)	(5)	(6)
% days low PBL height	0.586***	0.403***	0.374***	0.339***	0.403***	0.322***
	(0.152)	(0.124)	(0.118)	(0.118)	(0.122)	(0.110)
log 1806 population			0.005			0.006
			(0.008)			(0.008)
log 1836 central density			0.002			0.005
			(0.007)			(0.007)
% medium gleysoil depth				-0.273***		-0.280***
				(0.062)		(0.063)
Sq. % medium gleysoil depth				0.273***		0.281***
				(0.098)		(0.099)
% high organic carbon					-0.044	-0.027
					(0.058)	(0.052)
Sq. % high organic carbon					0.062	0.044
					(0.064)	(0.058)
Observations	2,232	2,232	2,214	2,232	2,232	2,214
Number of clusters (cities)	248	248	246	248	248	246
R^2	0.755	0.813	0.818	0.828	0.814	0.833
F-stat	112.60	63.87	21.43	83.66	24.07	37.99
Geography controls	X	X	X	X	X	X
Weather controls		X	X	X	X	X
Mean PBLH control		X	X	X	X	X
Year fixed effects	X	X	X	X	X	X

Notes: Geography controls include mean elevation, ruggedness, share of land covered by mountains, and dummies for sea access and waterway presence. Weather controls include summer temperature, winter temperature, mean wind speed and total precipitation.

D Additional reduced-form results and robustness

This appendix presents additional reduced-form results and robustness checks, in the following order:

- Additional results on the relationship between city structure and air pollution
- Robustness checks on the effect of local air pollution on nominal wage differentials, housing price differentials, and real wage differentials
- Additional results on the effect of local air pollution on working hours, total earnings, and value added per worker
- A replication of the main results at the commuting zone (CZ) level

D.1 Robustness: Local air pollution and wages

Table D.1: Impact of $PM_{2.5}$ level on the real wage "penalty"

	Outcome: (log) nominal wage W^{adj}					
	(1)	(2)	(3)	(4)	(5)	(6)
(log) density	0.026***	0.040***	0.014	0.031***	0.070**	0.044***
	(0.005)	(0.003)	(0.012)	(0.006)	(0.035)	(0.011)
(log) area		0.029***		0.030***		0.032***
		(0.002)		(0.002)		(0.002)
(log) $PM_{2.5}$					-0.675***	-0.129*
					(0.239)	(0.074)
Implied $\hat{\eta}_S = \hat{\beta}_S - \hat{\beta}_{N/S}$		-0.011		-0.001		-0.012
Density instrumented?			X	X	X	X
Area instrumented?				X		X
$PM_{2.5}$ instrumented?					X	X
Year fixed effects	X	X	X	X	X	X
Observations	2,241	2,241	2,241	2,232	2,223	2,223
Clusters (cities)	249	249	249	248	247	247
Mean outcome	0.523	0.523	0.523	0.523	0.523	0.523
F-stat, (log) density			30.57	41.68	38.26	33.60
F-stat, (log) area				367.16		299.69
F-stat, (log) $PM_{2.5}$					34.00	37.43
Sargan, p-value			0.001	0.003	0.001	0.154

Direct effect of local air pollution Table D.2 provides evidence of the direct effect of $PM_{2.5}$ on nominal wages in Columns 1-4, and on real wages, equal to nominal wages - (share of housing in expenditures $\times$ housing price index) in Columns 5-7. The results are consistent with the theoretical prediction of equations (36) and (38), as they show a consistently negative effect, apart from Columns 1-2, which show a positive correlation due to confounding factors.

Table D.2: Impact of local air pollution on nominal and real wages

	(log) nominal wage				*(log) real wage*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	OLS	TWFE	IV	IV	OLS	TWFE	IV
(log) $PM_{2.5}$	0.094***	0.002	-0.159*	-0.181	-0.301***	-0.055***	-0.136
	(0.021)	(0.004)	(0.082)	(0.116)	(0.050)	(0.012)	(0.115)
$PM_{2.5}$ instrumented?			X	X	X		X
Year fixed effects	X	X	X	X	X	X	X
City fixed effects		X				X	X
Weather controls	X	X	X	X	X	X	X
Geo controls	X		X	X	X		
Observations	4,628	4,628	4,615	2,295	2,295	2,295	2,295
F-stat, (log) $PM_{2.5}$			146	65			13

Notes: Standard errors clustered at the city (MSA) level in parentheses. $^{*}\ p < .10,\ ^{**}\ p < .05,\ ^{***}\ p < .01$
Geography controls include mean elevation, ruggedness, the share of city surface covered by mountains, and dummies for access to the sea and presence of a waterway. Weather controls include summer and winter average temperature, wind speed and total precipitation. IV for $PM_{2.5}$: number of days with an unusually low PBLH. Columns (1)-(4) use residual wages, controlling for individual sorting and sectoral composition, as an outcome. Columns (5)-(7) focus on real wages, i.e. residual wages of Columns (1)-(4) net of housing costs (share of housing in expenditure $\beta \times$ log housing price index) as an outcome. Column (4) restricts the sample to observations for which housing price data is available, so as to allow for a direct comparison with Columns (5)-(7).

Table D.3: Local air pollution and the urban wage premium: Robustness to alternative city size definitions

	Outcome: (log) nominal wage W^{adj}					
	(1)	(2)	(3)	(4)	(5)	(6)
	(a) City size definition: employment					
(log) city size	0.031***	0.040***	0.031***	0.032***	0.036***	0.044***
	(0.002)	(0.003)	(0.002)	(0.006)	(0.004)	(0.011)
(log) area		-0.011***		-0.0010		-0.013
		(0.003)		(0.007)		(0.011)
(log) $PM_{2.5}$					-0.014	-0.124*
					(0.051)	(0.070)
Observations	2,241	2,241	2,241	2,232	2,223	2,223
F-stat, (log) city size			221.09	225.93	37.190	179.70
F-stat, (log) area				246.44		210.94
F-stat, (log) $PM_{2.5}$					33.995	36.815
Sargan, p-value			0.001	0.008	0.079	0.175
	(b) City size definition: population density					
(log) density	0.026***	0.039***	0.004	0.033***	0.085**	0.047***
	(0.005)	(0.003)	(0.013)	(0.007)	(0.039)	(0.013)
(log) area		0.028***		0.030***		0.031***
		(0.002)		(0.002)		(0.003)
(log) $PM_{2.5}$					-0.733***	-0.126
					(0.253)	(0.082)
Observations	2,241	2,241	2,241	2,232	2,223	2,223
F-stat, (log) density			30.930	41.235	39.078	34.175
F-stat, (log) area				246.44		210.94
F-stat, (log) $PM_{2.5}$					33.995	36.815
Sargan, p-value			0.001	0.017	0.018	0.210
	(c) Including cities with missing housing price data					
(log) city employment	0.027***	0.035***	0.027***	0.038***	0.025***	0.045***
	(0.001)	(0.002)	(0.003)	(0.008)	(0.002)	(0.011)
(log) area		-0.009***		-0.012*		-0.018*
		(0.002)		(0.007)		(0.010)
(log) $PM_{2.5}$					0.013	-0.077
					(0.035)	(0.060)
Observations	6,069	6,069	6,069	6,001	4,615	4,615
F-test, (log) city size			70.630	1,701.8	1,106.9	1,107.5
F-test, (log) area				2,748.6		1,724.3
F-test, (log) $PM_{2.5}$					47.509	47.512
City size instrumented?			X	X	X	X
Area instrumented?			X	X	X	X
$PM_{2.5}$ instrumented?					X	X
Year fixed effects	X	X	X	X	X	X

Notes: Standard errors clustered at the city (MSA) level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$
All panels of this Table replicate Table 3, such that all regressions include for geography and weather controls. Instruments for city size and surface area are soil characteristics and historical population. Instrument for $PM_{2.5}$ is the number of days with a low PBL height. When $PM_{2.5}$ is instrumented, regressions include a control for city mean PBL height over the study period.

D.2 Validity checks: air pollution and working hours, earnings and value added per worker

Table D.4: Elasticity of working hours and labour earnings to density and air pollution

	(1)	(2)	(3)	(4)	(5)	(6)
	(a) Outcome: (log) raw number of working hours					
(log) density	-0.023***	-0.049**	0.029	-0.130*	0.034	-0.158**
	(0.006)	(0.018)	(0.035)	(0.069)	(0.030)	(0.077)
(log) area		0.031		0.103		0.134*
		(0.020)		(0.070)		(0.074)
(log) $PM_{2.5}$					-0.356	0.112
					(0.229)	(0.405)
Observations	2,155	2,155	2,155	2,137	2,146	2,128
F-stat, (log) city size			37.358	239.19	39.936	213.50
F-stat, (log) area				258.22		238.49
F-stat, (log) $PM_{2.5}$					58.938	47.665
Sargan, p-value			0.0001	0.0001	0.0001	0.001
	(b) Outcome: (log) adjusted number of working hours					
(log) density	-0.097***	-0.061**	-0.154**	-0.338	-0.182***	-0.289*
	(0.024)	(0.027)	(0.052)	(0.218)	(0.051)	(0.162)
(log) area				0.164		0.116
				(0.211)		(0.154)
(log) $PM_{2.5}$					-0.543	-0.043
					(0.383)	(0.724)
Observations	2,155	2,155	2,155	2,137	2,146	2,128
F-stat, (log) density			37.358	239.19	39.936	205.82
F-stat, (log) area				258.22		216.78
F-stat, (log) $PM_{2.5}$					58.938	44.503
Sargan, p-value			0.001	0.001	0.001	0.001
	(c) Outcome: (log) adjusted labour earnings					
(log) density	-0.090***	-0.048	-0.139**	-0.335	-0.169***	-0.274*
	(0.025)	(0.027)	(0.050)	(0.213)	(0.050)	(0.157)
(log) area		-0.050		0.166		0.108
		(0.047)		(0.206)		(0.149)
(log) $PM_{2.5}$					-0.537	-0.093
					(0.381)	(0.695)
Observations	2,155	2,155	2,155	2,137	2,146	2,128
F-test, (log) city size			37.358	239.19	39.936	204.62
F-test, (log) area				258.22		216.47
F-test, (log) $PM_{2.5}$					58.938	44.185
City size instrumented?			X	X	X	X
Area instrumented?			X	X	X	X
$PM_{2.5}$ instrumented?					X	X
Year fixed effects	X	X	X	X	X	X

Sources: INSEE, DADS, IGN, ERA5, ESDB.

Notes: Standard errors clustered at the city (MSA) level in parentheses. $^{*}\ p < .10,\ ^{**}\ p < .05,\ ^{***}\ p < .01$

All panels of this Table replicate Table 3, such that all regressions include for geography and weather controls. Instruments for city size and surface area are soil characteristics and historical population. Instrument for $PM_{2.5}$ is the number of days with a low PBL height. When $PM_{2.5}$ is instrumented, regressions include a control for city mean PBL height over the study period.

Outcomes are the following: in Panel (a), the raw average number of working hours at the city×year level, in Panel (b), the average number of working hours adjusted for sectoral composition and worker fixed effects in the fashion of equation (39) for wages, in Panel (c) average labour earnings, adjusted in the same fashion.

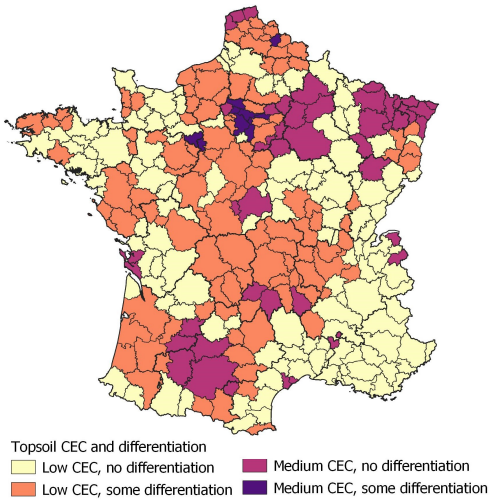
D.3 Robustness: A replication at the commuting zone level

In the main text, I use on INSEE's definition of urban areas equivalent to U.S. Metropolitan Statistical Areas (MSAs), close the European Union's Functional Urban Areas (FAUs). This appendix shows the robustness of the main results to a change in the city definition, by looking at commuting zones (CZ). The way each delineation is built is discussed in Section 3.6. Figure D.1 below displays the spatial delineation of France into its 297 commuting zones. As one can see, the Paris region is split into several commuting zones, in a way that is forced by INSEE to prevent it from getting too large. This is part of the reason why I focus on the MSA definition in the main text: continuities in the urban fabric and commuting flow intensity are preserved by the MSA delineation, while the CZ delineation breaks these rules.

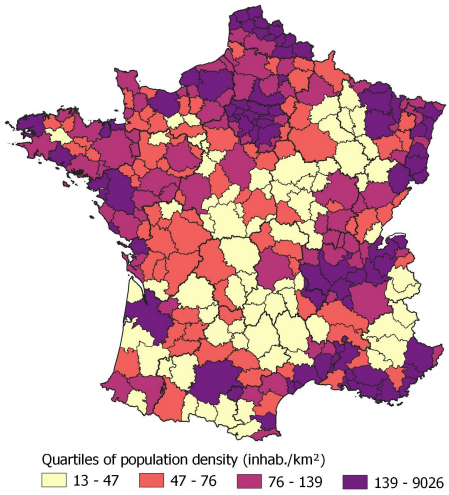
The summary statistics of all variables are shown at the CZ level in Table D.5 below.

Figure D.1: Instrumental variables (left) and endogenous regressors (right) at the CZ level

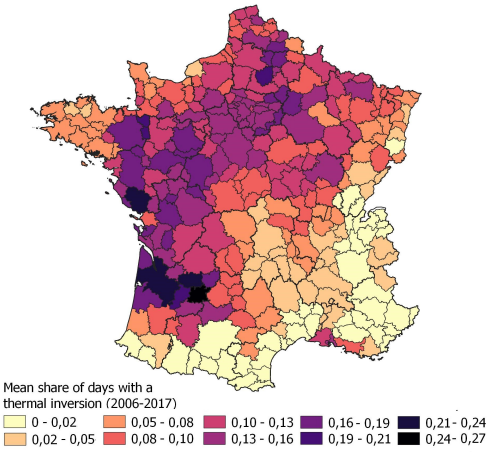
(a) Soil characteristics



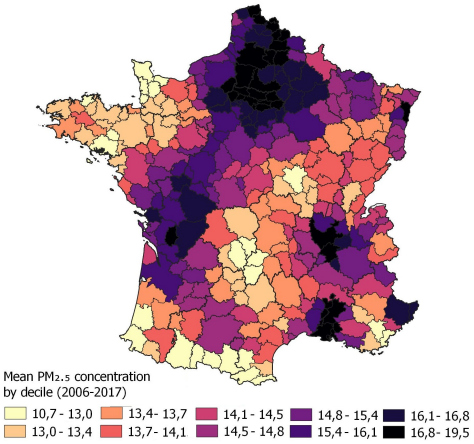
(b) 2010 population density



(c) % of days with a low PBLH



(d) Mean $PM_{2.5}$ concentration



Note: Author's elaboration based on INSEE, IGN, Copernicus ERA5 and ACAG data.

Table D.5: Summary statistics - French commuting zones, 2002-2018

Variable	Mean	Std Dev.	Min.	Max.	N×T
Main variables					
Hourly log wage W^{raw} (2019 €)	2.40	0.06	2.26	2.67	5,049
Adjusted log wage W^{adj}	-7.54	2.76	-12.01	-3.07	5,049
Adjusted log wage W^{adj}, demeaned	0	0.027	-0.08	0.11	5,049
Employment density	73.91	327.5	3.77	5,392	5,049
Log employment density	3.48	0.94	1.33	8.59	5,049
Population density	178.6	577.0	10.52	9,139	5,049
Log population density	4.47	0.94	2.35	9.12	5,049
Population (thousands)	211.3	410.7	8.42	6,002	5,049
Log population	11.70	0.95	9.04	15.61	5,049
Area (sq. km)	1,819	1,396	119.2	8,826	5,049
Log area	7.24	0.75	4.78	9.09	5,049
$PM_{2.5}$ concentration ($\mu g/m^3$)	12.61	2.11	7.89	19.51	5,049
Log PM concentration	2.52	0.17	2.07	2.97	5,049
Number of $PM_{2.5}$ pixels	2,136	1,617	148	9,824	5,049
Log housing price (€/km^2)	-0.30	0.37	-1.34	0.71	1,405
Weather controls					
Mean winter temperature (°C)	6.63	2.11	-6.75	11.88	5,049
Mean summer temperature (°C)	16.39	1.75	7.25	22.15	5,049
Mean wind speed (km/h)	12.49	3.51	4.26	26.81	5,049
Mean cloud cover (%)	61.97	6.42	35.94	73.85	5,049
Total precipitation (m)	0.83	0.23	0.30	2.01	5,049
Geographical controls					
Mean elevation (m)	300.83	343.63	0.22	2,286	5,049
Ruggedness (m)	379.34	526.92	15	4,030	5,049
Sea (dummy, %)	18.18				5,049
Waterway (dummy, %)	50.17				5,049
Mountainous (dummy, %)	33.46				5,049
Instruments					
Log 1806 population	10.50	0.88	7.58	13.44	5,049
Topsoil CEC: medium (%)	17.8				5,049
Some soil differentiation (%)	41.8				5,049
Number of days with low PBLH	18.15	19.18	0	160	5,049

Notes: Author's computation based on DADS, INSEE, ACAG, the Copernicus project, IGN, Observatoire des Territoires, Cassini-EHESS, DV3F and ESDB information.

Table D.6: Impact of density on air pollution at the CZ level

	Outcome: (log) $PM_{2.5}$					
	(1) OLS1	(2) OLS2	(3) TSLS1	(4) TSLS2	(5) Pop.	(6) Subs.
(log) density	0.047***	0.043***	0.065***	0.062***	0.067***	0.060***
	(0.003)	(0.004)	(0.010)	(0.008)	(0.009)	(0.009)
(log) area		-0.014***		-0.009	-0.053***	-0.008
		(0.005)		(0.007)	(0.013)	(0.008)
F-stat, (log) density			346.7	425.9	3,117	112.6
F-stat, (log) area				1,281	1,281	335.8
Sargan p-value			0.036	0.161	0.816	0.390
Year fixed effect	X	X	X	X	X	X
Geographical controls	X	X	X	X	X	X
Weather controls	X	X	X	X	X	X
Observations	5,049	5,049	5,049	5,049	5,049	1,405
Cities (clusters)	297	297	297	297	297	281

Note: Standard errors clustered at the city (CZ) level in parentheses. $^{*}\ p < .10,\ ^{**}\ p < .05,\ ^{***}\ p < .01$
Cities are defined as employment zones (commuting zones). Geographical controls include mean elevation, ruggedness, and dummies for access to the sea, presence of a waterway, and mountains. Weather controls include summer and winter average temperature, wind speed and cloud cover. Columns 3-6 use topsoil CEC and differentiation as instruments for employment density. Columns 4-6 use 1806 population as an instrument for land area. Column 6 focuses on the subsample for which housing prices are available.

Table D.7: Impact of density on air pollution - Clay instruments - CZ level

	Outcome: (log) $PM_{2.5}$				
	(1) OLS	(2) OLS	(3) TSLS	(4) TSLS	(5) Pop.
(log) density	0.0471***	0.0432***	0.0727***	0.0700***	0.0818***
	(0.0033)	(0.0036)	(0.0149)	(0.0135)	(0.0160)
(log) area		-0.0134***		-0.0042	-0.0816***
		(0.0051)		(0.0119)	(0.0192)
Year fixed effect	X	X	X	X	X
Geographical controls	X	X	X	X	X
Weather controls	X	X	X	X	X
Observations	5,049	5,049	5,049	5,049	5,049
Cities (clusters)	297	297	297	297	297
Adjusted R^2	0.774	0.777	0.774	0.776	0.781
F-stat, (log) density			346.7	425.9	3,117
F-stat, (log) area				1,281	1,281
Sargan p-value			0.036	0.161	0.816

Note: Standard errors clustered at the city (CZ) level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$
Cities are commuting zones. Geographical controls include mean elevation, ruggedness, and dummies for access to the sea, presence of a waterway, and mountains. Weather controls include summer and winter average temperature, wind speed and cloud cover. Columns 3-6 use clay instruments as instruments for employment density. Columns 4-5 use 1806 population as an instrument for land area. Column 5 focuses on the subsample for which housing prices are available.

Table D.8: Impact of density on emissions and commuting

	log residential emissions (1)	log traffic emissions (2)	log traffic em. /worker (3)	log traffic em. /car commuter (4)	log dist. travelled (5)	log dist. /car commuter (6)
(log) density	0.520***	0.472***	-0.529***	-0.411***	0.859***	-0.024**
	(0.047)	(0.039)	(0.039)	(0.043)	(0.020)	(0.010)
(log) area	-0.030	-0.073**	-1.073***	-1.079***	1.336***	0.332***
	(0.041)	(0.036)	(0.036)	(0.034)	(0.013)	(0.008)
Implied $\hat{\phi}_S = \hat{\beta}_S - \hat{\beta}_{N/S}$	-0.550	-0.545	-0.544	-0.668	0.477	0.356
Year fixed effects	X	X	X	X	X	X
Geographical controls	X	X	X	X	X	X
Weather controls	X	X	X	X	X	X
Observations	5,049	5,049	5,049	3,861	3,861	3,861
Cities	297	297	297	297	297	297
F-stat, (log) density	433.98	433.98	433.98	331.75	331.75	331.75
F-stat, (log) area	1,262.6	1,262.6	1,262.6	962.07	962.07	962.07
Cragg-Donald	225.38	225.38	225.38	170.26	170.26	170.26
Wu-Hausman, p-value	10^{-22}	10^{-10}	10^{-10}	10^{-8}	0.004	10^{-35}
Sargan, p-value	0.003	0.002	0.002	0.074	0.502	0.001

Notes: Standard errors clustered at the city (CZ) level in parentheses. $^{*}\ p < .10,\ ^{**}\ p < .05,\ ^{***}\ p < .01$
Data on emissions from EDGAR. Number of car commuters built using the restricted access version of the Census. Each car commuter is assumed to commute from the population-weighted centroid of the municipality they reside in to the population-weighted centroid of the municipality they work in, *via* the shortest path by road, including tolls.

Table D.9: Impact of density on nominal and real wage differentials – CZ level

	(log) wage W^{adj}			*(log) real wage* W^R	
	(1)	(2)	(3)	(4)	(5)
(log) density	0.020***	0.018***	0.024***	-0.073***	-0.046**
	(0.001)	(0.005)	(0.009)	(0.009)	(0.022)
(log) $PM_{2.5}$			-0.223**		-0.755**
			(0.108)		(0.011)
Density IV?		X	X	X	X
$PM_{2.5}$ IV?			X		X
Year fixed effects	X	X	X	X	X
Geo/weather controls	X	X	X	X	X
Mean PBLH control			X		X
Observations	5,049	5,049	5,049	1,405	1,405
Cities (clusters)	297	297	297	281	281
F-stat (1st stage), (log) density		140.88	163.26	112.6	98.1
F-stat (1st stage), (log) area		403.16	345.05	349.1	248.6
F-stat (1st stage), (log) $PM_{2.5}$			81.53		25.4
Sargan, p-value		0.268	0.218		

Note: Standard errors clustered at the city (CZ) level in parentheses. $^{*}\ p < .10,\ ^{**}\ p < .05,\ ^{***}\ p < .01$
IV for city density: soil characteristics. IV for city surface area: 1806 population. IV for $PM_{2.5}$: number of days with an unusually low PBLH.

Table D.10: Impact of air pollution on nominal and real wages – CZ level

	(log) wage W^{adj}			*(log) real wage W^R*	
	(1)	(2)	(3)	(4)	(5)
(log) $PM_{2.5}$	0.071***	-0.178**	-0.223**	0.044***	-0.096*
	0.011	0.083	0.108	0.019	0.058
Density IV?			X		
$PM_{2.5}$ IV?		X	X		X
Year fixed effects	X	X	X	X	X
Geo/weather controls	X	X	X	X	X
Mean PBLH control		X	X		X
Observations	5,049	5,049	5,049	5,049	5,049
Cities (clusters)	297	297	297	297	297
F-stat (1st stage), (log) area		403.16	345.05		403.16
F-stat (1st stage), (log) $PM_{2.5}$		44.35	81.53		44.35
Sargan, p-value			0.218		

Notes: Standard errors clustered at the city (CZ) level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$
IV for density: soil characteristics. IV for area: 1806 population. IV for $PM_{2.5}$: number of days with an unusually low PBLH.

Table D.11: Impact of local air pollution on (raw) value added per worker

	(1)	(2)	(3)	(4)	(5)
			Panel 1: Logarithms		
(log) PM	0.003	-0.396	-0.235**	-0.202**	-0.225
	(0.025)	(0.727)	(0.083)	(0.077)	(0.217)
(log) density			0.084***	0.086***	0.062***
			(0.004)	(0.003)	(0.019)
(log) area				0.008**	-0.008
				(0.004)	(0.010)
			Panel 2: Levels		
$PM_{2.5}$	-0.0006	-0.032	-0.026**	-0.015**	-0.036***
	(0.0009)	(0.103)	(0.010)	(0.006)	(0.013)
(log) density			0.093***	0.087***	0.098***
			(0.006)	(0.004)	(0.008)
(log) area				0.009**	-0.0006
				(0.004)	(0.012)
F-stat (1st stage), (log) PM		7.37	187.63	383.69	114.16
F-stat (1st stage), (log) density			272.84	1,086.6	225.96
F-stat (1st stage), (log) area				167.41	390.83
Kleibergan-Paap		10.04	32.45	22.05	17.21
$PM_{2.5}$ instrumented?		X	X	X	X
Mean PBLH control			X	X	X
Year fixed effects	X	X	X	X	X
City fixed effects	X	X			
Observations	5,049	5,049	5,049	5,049	5,049

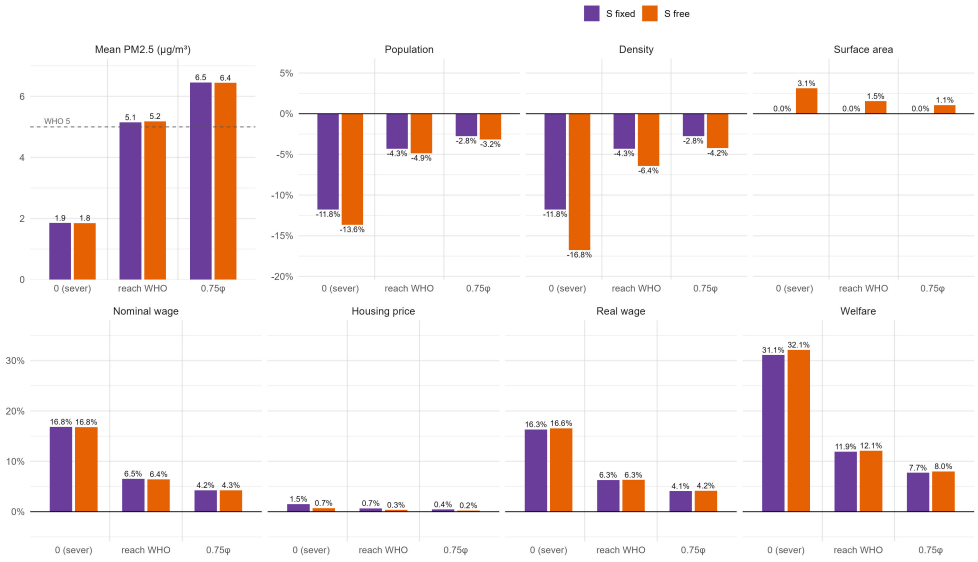
Note: Standard errors clustered at the city (MSA) level in parentheses. $^{*}\ p < .10,\ ^{**}\ p < .05,\ ^{***}\ p < .01$
All regressions control for summer and winter average temperature, wind speed, cloud cover and total precipitation.
IV for $PM_{2.5}$: number of days with an unusually low PBL height.

E Additional structural results and robustness

E.1 Bringing cities to the WHO guideline

Figure E.1: Meeting the WHO guideline: city-level and baseline population results

(a) City-level outcomes (no weight)



(b) Weighted by baseline population

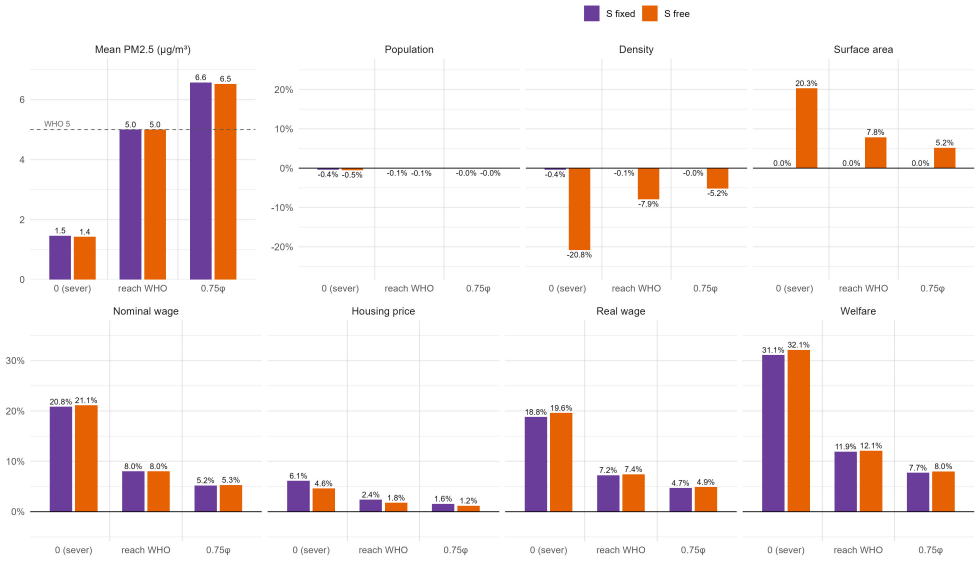


Figure E.2: Meeting the WHO guideline: channel decomposition at the worker level



Figure E.3: Meeting the WHO guideline: channel decomposition at the city level



E.2 Projecting current internal migration trends

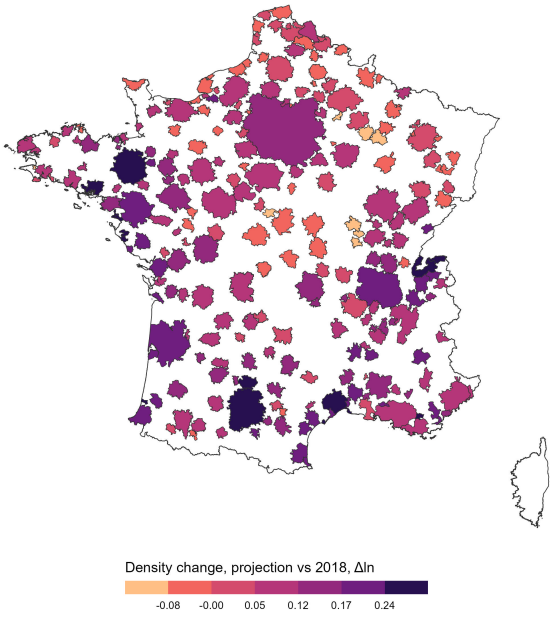
I provide the details behind the computation of the third counterfactual exercises, which extrapolates current city population changes.

Backing out amenity shock θ I proceed in two steps. I start from the 2018 equilibrium, and recover each city's fundamentals (productivity, amenities, construction productivity and pollution endowment) from observed population, wages, housing prices, $PM_{2.5}$ concentration and surface area, as in Section 6. I then set a target population for each city, assuming that its population grows over the next nineteen years by the same proportion as between 1999 and 2018, as shown in Panel (E.4a) of Figure E.4 below. For each city, I back out the level of amenities that delivers exactly this target population, holding all other fundamentals and air pollution at their 2018 levels. I pin down the common utility level by requiring these amenity shocks to average to zero across cities, weighted by target population. The exercise therefore reallocates amenities across cities rather than making all cities more attractive, and the growth of national population is absorbed by a change in common utility. This backed-out change in perceived amenities in shown in the map of Panel (E.4b).

Projecting the 1999-2018 trend In the second step, I apply these amenity shocks to the 2018 fundamentals and solve the model again, this time letting air pollution respond to the new population and surface area through equation (15). Because growing cities also become more polluted, they grow by less than their target, and comparing this outcome with the one where air pollution is held at its 2018 level isolates the role of the pollution feedback. I repeat this second step with both pollution elasticities ϕ_N and ϕ_S halved, to see how a partial decoupling of air pollution from city structure changes the results. I run the whole exercise with a fixed and with an endogenous city surface (in the latter case, the surface also adjusts in both steps). All outcomes are reported as changes relative to the 2018 equilibrium.

Figure E.4: 1999-2018 population change and deduced perceived amenity change

(a) Observed population change



(b) Backed-out amenity change

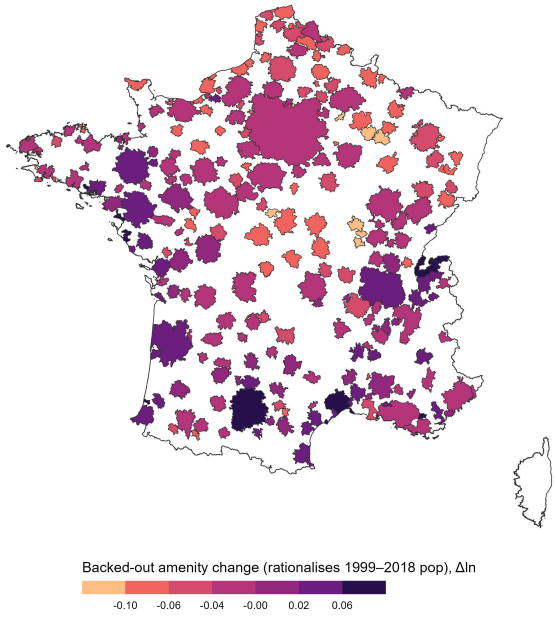
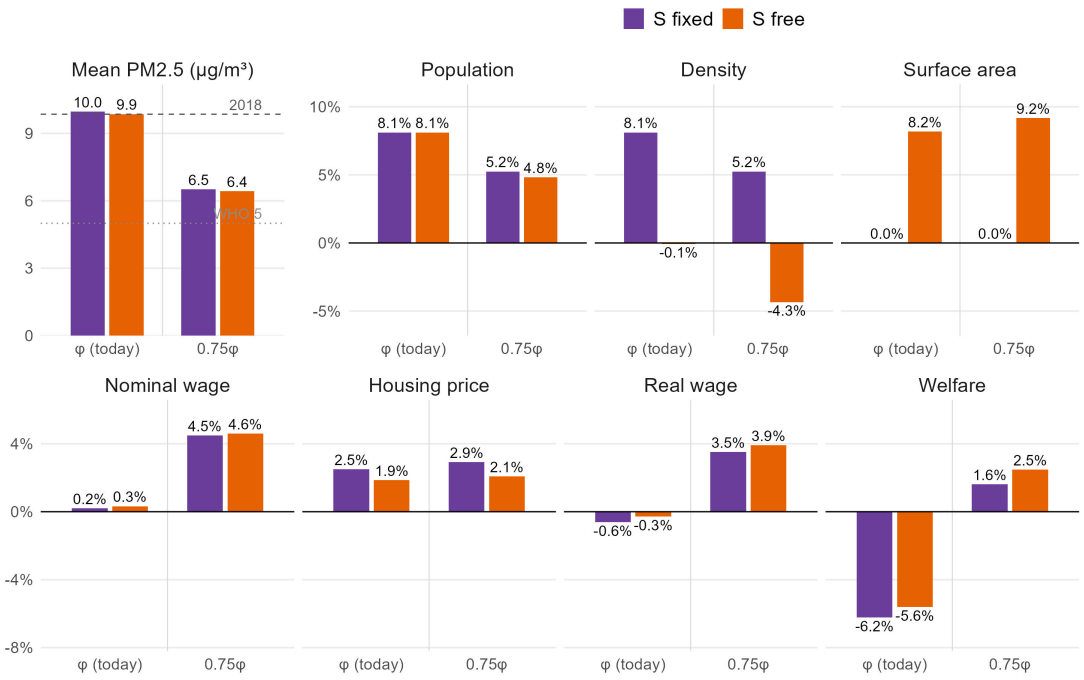


Figure E.5: Projecting migration trends: city-level and baseline population results

(a) City-level outcomes (no weight)



(b) Weighted by baseline population

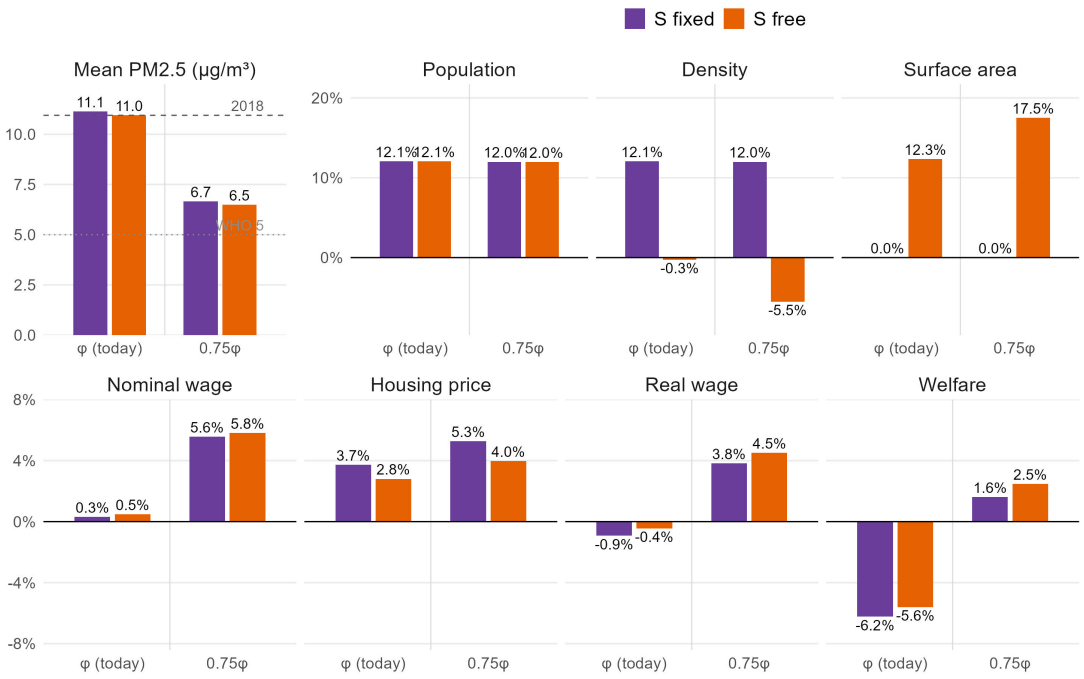
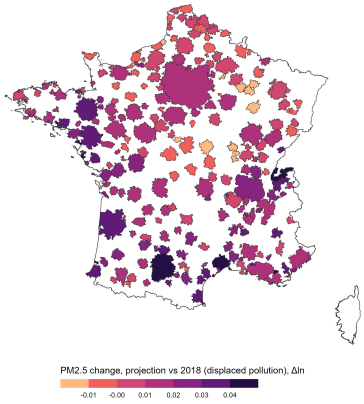
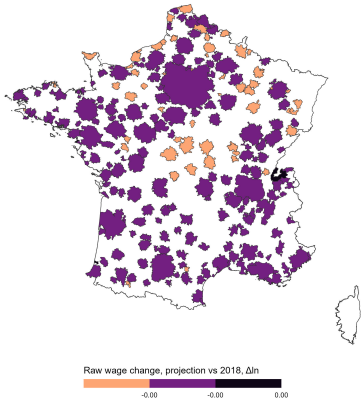


Figure E.6: Projecting migration trends: effects city by city

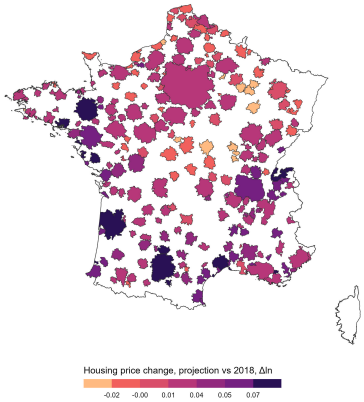
(a) Air pollution



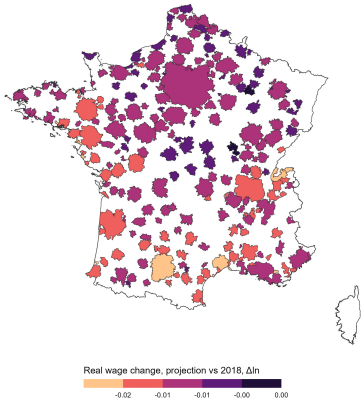
(b) Raw wage



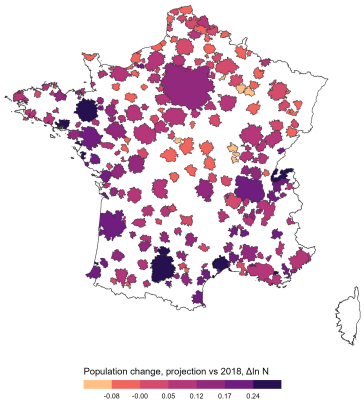
(c) Housing prices



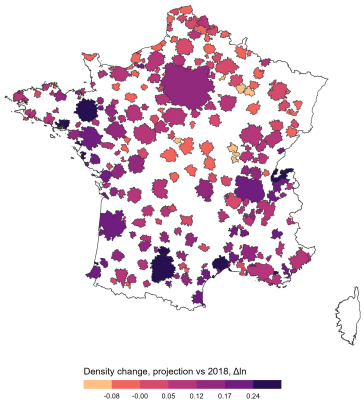
(d) Real wage



(e) Population



(f) Density



Notes: City surface area is held fixed.