

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Migration and hedonic valuation: The case of air quality

Patrick Bayer^{a,b}, Nathaniel Keohane^c, Christopher Timmins^{a,b,*}^a Department of Economics, Duke University, P.O. Box 90097, Durham, NC 27708, USA^b NBER, Cambridge, MA, USA^c Environmental Defense Fund, 257 park Avenue S., New York, NY 10010

ARTICLE INFO

Article history:

Received 24 January 2008

Available online 23 January 2009

Keywords:

Particulate matter

Valuation of air quality

Wage-hedonic models

Migration costs

Residential sorting

Discrete-choice models

ABSTRACT

Conventional hedonic techniques for estimating the value of local amenities rely on the assumption that households move freely among locations. We show that when moving is costly, the variation in housing prices and wages across locations may no longer reflect the value of differences in local amenities. We develop an alternative discrete-choice approach that models the household location decision directly, and we apply it to the case of air quality in US metro areas in 1990 and 2000. Because air pollution is likely to be correlated with unobservable local characteristics such as economic activity, we instrument for air quality using the contribution of distant sources to local pollution—excluding emissions from local sources, which are most likely to be correlated with local conditions. Our model yields an estimated elasticity of willingness to pay with respect to air quality of 0.34–0.42. These estimates imply that the median household would pay \$149–\$185 (in constant 1982–1984 dollars) for a one-unit reduction in average ambient concentrations of particulate matter. These estimates are three times greater than the marginal willingness to pay estimated by a conventional hedonic model using the same data. Our results are robust to a range of covariates, instrumenting strategies, and functional form assumptions. The findings also confirm the importance of instrumenting for local air pollution.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Since Rosen's [13] seminal paper, economists have used hedonic techniques to estimate the value of a wide range of amenities, including clean air, school quality, and lower crime rates. The great attraction of the approach is that it uses observed behavior in housing and labor markets to infer the value of non-market goods. On the standard assumption that individuals choose the residential locations that maximize their utility, marginal rates of substitution between local amenities and other goods will equal the price ratio. Hence the marginal willingness to pay (MWTP) for those amenities can be measured by their implicit prices, as reflected in housing prices and wages. The broad avail of this approach, along with considerable practical interest in the estimates it provides, explains the continuing interest among economists in the theory and identification of hedonic models. The topic has, moreover, recently seen a resurgence with methodological innovations [3,8].

This paper addresses a crucial but often overlooked assumption in hedonic models, and shows how that assumption may lead to biased estimates of MWTP for local amenities. Hedonic models typically assume that people can move freely among locations when they buy homes and choose jobs. If so, wages and rents must adjust to reflect the implicit prices of

* Corresponding author at: Department of Economics, Duke University, P.O. Box 90097, Durham, NC 27708, USA. Fax: +1919 684 8974.
E-mail address: christopher.timmins@duke.edu (C. Timmins).

local amenities; hence, MWTP can be inferred from variation in housing prices and income. In reality, migration is costly; moving to a new city entails both out-of-pocket costs and the psychological costs of leaving behind one's family and cultural roots. Data on residential choices suggest that such costs are significant. Table 1 relates birth location to residential location; it shows that the great majority of US household heads reside in the region of their birth. A similar pattern holds at the state level. This strong revealed preference for staying close to home belies the assumption that residential choices reflect a simple tradeoff between local attributes and prevailing rents and wages. If migration costs enter into residential location decisions, they should be considered by analysts measuring the value of local amenities.

How will migration costs affect estimates of MWTP? Consider an exogenous improvement in air quality in a particular city. In response, we would expect housing prices to rise and wages to fall until a new equilibrium is reached. If migration is costless, these changes will fully reflect the value of the cleaner air. But if migration is costly, the change in housing prices and wages must be smaller; the benefit people get from moving to the city must now compensate them not only for the higher rents and lower income, but also for the cost of moving. MWTP estimates that ignore these moving costs will be biased.

Beyond the theoretical questions of identification and estimation, numerical estimates of the value of local public goods are of great practical interest. Again consider the example of air quality, whose protection motivates a range of government policies that impose substantial costs on firms and consumers. A comprehensive survey of cross-sectional hedonic property value studies found wide dispersion in estimated willingness to pay, with many instances of negligible or even negative estimates [18]. If those low estimates are reliable, the costs of stringent air pollution regulation may outweigh the benefits. On the other hand, evidence that such estimates understate the value of clean air would bolster the case for government policy.

In this paper, we show how migration costs can be incorporated into a hedonic analysis. We start by incorporating migration into the canonical wage-hedonic model [12]. If moving is costly, then the sum of the derivatives of housing prices and wages with respect to the amenity – the standard hedonic measure of MWTP – will no longer equal the implicit price of the amenity. The costlier the migration relative to the marginal benefits of an improvement in the amenity, the greater the bias from ignoring migration costs in the analysis. To allow for costly mobility, we employ a different empirical strategy. The starting point for our analysis is the household location decision, rather than the first-order condition implied by a traditional hedonic model. This approach allows us to incorporate migration costs (as the implicit disutility of moving various distances from one's birth state) directly into the household optimization problem.

We apply our method to the case of air quality—specifically, ambient concentrations of particular matter (“PM10”) in metropolitan areas throughout the US, for the years 1990 and 2000. We study air pollution in general, and PM10 in particular, for a number of reasons. First, an estimate of the economic value of improvements in air quality is of central importance to the US Environmental Protection Agency (EPA) in the regulation of air pollution under the Clean Air Act and its subsequent amendments. Second, air quality improved significantly over the decade studied, providing useful panel variation. Third, migration costs are likely to be large relative to the potential gains of changing locations for the sake of air quality; hence ignoring such costs is likely to produce substantial bias in estimates of MWTP for air quality.¹ Fourth, a long literature has used hedonic methods to value air quality [10,14,18]. Finally, particulate matter is a natural choice of pollutant. It is the standard measure of air pollution used in the literature, and an increasing body of evidence suggests that it is by far the most important local air pollutant in terms of health effects.

Our empirical analysis proceeds in two stages. First, we use a discrete-choice model to infer the utility associated with living in various metropolitan areas. We then regress these metro-area utilities on air pollution concentrations in order to recover the MWTP for air quality. This second stage is analogous to the traditional hedonic approach, which regresses housing prices on air pollution. An identification problem thus arises that is endemic to hedonic analyses. Local air quality is likely to be correlated with unobserved local economic factors, which also affect housing prices. If so, naïve estimates of willingness to pay will be biased downward—helping to explain the low estimates reported in the existing literature.

We employ a novel instrumental variables (IV) approach to deal with this endogeneity problem. The intuition behind our approach is simple. Although local emissions (correlated with local economic activity) are the major determinant of local air quality, pollution also wafts in from distant sources. The tall stacks of electric power plants spew particulate matter and other pollutants high into the atmosphere, where they travel great distances before affecting ground-level air quality. Distant emissions, however, are likely to be uncorrelated with local economic activity—a conjecture that is confirmed by the data. Hence pollution from distant sources provides a natural instrument for local air pollution. We compute this instrument using a detailed source-receptor (S-R) matrix developed for the US EPA that relates emissions from nearly 6000 sources to particulate matter concentrations in each county in the US.

Our results demonstrate the importance of accounting for endogeneity and incorporating mobility costs. As a preliminary step, we estimate a traditional wage-hedonic model. Instrumenting for air pollution is found to greatly increase the magnitude of the estimated coefficient on particulate matter concentration in a regression of housing prices on local amenities. These initial results provide a benchmark for assessing the results of our residential sorting model. In line

¹ As a likely contrast, consider the case of households sorting across school districts within a single metropolitan statistical area in response to changes in school quality. Here we would expect that migration costs would be low, and that households would be highly motivated, leading to an expectation that the bias may be quite small.

Table 1
Regional mobility patterns—percent birth region by residence region as of the year 2000 (US Census data).

| Birth region | Residence in 2000 | | | | | | | | |
|--------------------|-------------------|--------------|--------------------|--------------------|----------------|--------------------|--------------------|----------|---------|
| | New England | Mid-Atlantic | East North Central | West North Central | South Atlantic | East South Central | West South Central | Mountain | Pacific |
| New England | 65.02 | 5.87 | 2.35 | 0.94 | 12.68 | 0.47 | 1.88 | 3.05 | 7.75 |
| Mid-Atlantic | 4.03 | 63.34 | 4.55 | 1.03 | 18.04 | 0.22 | 2.13 | 2.49 | 4.18 |
| East North Central | 0.59 | 1.60 | 73.42 | 2.83 | 9.84 | 2.09 | 2.78 | 2.94 | 3.90 |
| West North Central | 0.36 | 1.77 | 7.80 | 57.62 | 7.27 | 1.24 | 5.85 | 8.51 | 9.57 |
| South Atlantic | 0.99 | 3.59 | 4.50 | 0.84 | 79.47 | 2.82 | 2.67 | 1.60 | 3.51 |
| East South Central | 1.72 | 1.29 | 7.08 | 0.86 | 15.45 | 63.73 | 4.94 | 1.29 | 3.65 |
| West South Central | 0.47 | 1.54 | 2.13 | 1.42 | 6.51 | 2.49 | 77.75 | 3.31 | 4.38 |
| Mountain | 0.89 | 1.11 | 3.54 | 2.43 | 3.54 | 1.11 | 5.09 | 69.03 | 13.27 |
| Pacific | 0.86 | 1.61 | 3.42 | 1.52 | 5.13 | 1.14 | 2.75 | 7.69 | 75.88 |

Notes: Rows indicate birth regions; columns denote current residence. For example, the upper-left-hand cell indicates that 65.02% of household heads born in New England were living in the region during the 2000 Census. Regions are assigned according to Census definitions: Regional Definitions: (1) *New England* (CT, ME, MA, NH, RI, VT), (2) *Middle Atlantic* (NJ, NY, PA), (3) *East North Central* (IL, IN, MI, OH, WI), (4) *West North Central* (IA, KS, MN, MO, NE, SD, ND), (5) *South Atlantic* (DE, DC, FL, GA, MD, NC, SC, VA, WV), (6) *East South Central* (AL, KY, MS, TN), (7) *West South Central* (AR, LA, OK, TX), (8) *Mountain* (AZ, CO, ID, MT, NV, NM, UT), and (9) *Pacific* (AK, CA, HI, OR, WA).

with intuition, the estimated value of clean air rises considerably (by a factor of three, compared to the standard hedonic methodology) when migration costs are taken into account. Importantly, we show that these results are robust to a range of alternative specifications. These results have important implications for policy, suggesting that the economic benefits of regulations that reduced particulate matter emissions are substantially larger than found in previous studies that have ignored migration costs.

2. Econometric models for valuing local amenities

2.1. Incorporating mobility costs into the traditional hedonic model

Consider the following variant of Roback’s [12] model, incorporating mobility costs. We present the simplest possible version of this model in order to demonstrate the basic intuition. At the end of the section, we argue that extending the model to make it more realistic will only exacerbate the difficulties introduced by mobility costs.

As in Roback’s model, all individuals simultaneously choose their location along with consumption of a composite commodity C and a non-traded good (“housing”) H . Each location j is characterized by a quantity X_j of a location-specific amenity (“air quality”). In addition, there is a moving cost M_j associated with settling in city j . We treat M_j as a long-run migration cost (i.e., the cost incurred by adults choosing where to live relative to their birthplace). As such, these costs are primarily psychological and do not appear in the budget constraint. Following Roback, we assume that individuals have identical preferences and abilities. To keep the model as simple as possible, we suppose that all individuals are born in the same place, and that moving costs are a monotonic function of the amenity level. For example, we might imagine that everyone is born in a central location, and that other cities are arranged in concentric rings with amenities improving as one moves outward.

Individuals choose their location j , along with consumption of C and H , to maximize their utility subject to a budget constraint:

$$\max_{(C,H,X_j)} U(C,H;X_j,M_j) \text{ s.t. } C + \rho_j H = I_j, \tag{1}$$

where I_j is income in location j ; ρ_j is the price of housing in location j ; and the price of the composite commodity is normalized to unity. In equilibrium, individuals must be indifferent among locations; if not, they would prefer to move. Hence indirect utility, denoted \bar{V} , is constant; $V(I_j, \rho_j; X_j, M_j) \equiv \bar{V}$.

Individuals trade off local amenities against wages and rents (which affect the budget constraint and determine their consumption of commodity C). Taking the total derivative of indirect utility and using Roy’s Identity to substitute for $H = -V_p/V_r$, we arrive at the following equation for the implicit amenity price p^* :

$$p^* = H \frac{d\rho}{dX} - \frac{dI}{dX} - \frac{V_M}{V_r} \frac{dM}{dX}. \tag{2}$$

Hence p^* is the MWTP—more precisely, the change in income that would exactly compensate the individual for a marginal change in the amenity at their chosen location. The first two terms on the right-hand side of Eq. (2) are the familiar terms

from Roback's analysis. If mobility is costless ($V_M = 0$), or mobility costs are constant ($dM = 0$), then the model is identical to Roback's. In those cases, the implicit price of the amenity X can be measured as the extra cost of housing minus the compensating wage increase.

When mobility costs are positive and vary with location, the familiar equation no longer holds. Suppose that the amenity increases with distance from 0. In this case, $V_M < 0$ (since mobility is costly) and $dM/dX > 0$. Thus the true value of a marginal change in the amenity, given by p^* , is greater than the sum of the housing price and wage effects. Intuitively, when it is more costly to move to locations with better amenities, the housing and labor markets will appear to undervalue those amenities. In order to induce anyone to move to the more attractive locales, rents must be lower (or wages higher) than they would in a world without mobility costs.

Even this simple model poses difficulties for empirical analysis. If moving costs could be directly observed, then dM/dX could be estimated much as the housing-price ($d\rho/dX$) and income (dI/dX) gradients are, and the implicit price p^* could be inferred. But M is likely to be unobservable, since it represents the disutility of moving to an unfamiliar place far from home. Moreover, the restrictive assumptions we have made so far amount to the best-case scenario for the traditional model. For example, suppose that individuals are born in different locations. Then the requirement that all individuals' utilities equal the same constant \bar{V} no longer holds, invalidating the total differential approach, which determines the key marginal conditions given by Eq. (2). Or suppose that mobility costs do not vary systematically with location; then there is no longer any reason to expect that the implicit price must be equal across locations, which is the central identifying assumption of the hedonic model. We conclude that when mobility costs are likely to be significant, a different empirical strategy is necessary.

2.2. A model of residential sorting

To surmount these difficulties, we develop a structural approach that explicitly models the location decision as taking place prior to the consumption of housing and the composite commodity. Essentially, we push the analysis back a step, examining the utility maximization problem in (1) rather than simply analyzing the equilibrium condition implicit in (2). Estimation proceeds in two steps. First, we specify a discrete-choice model of the household location decision. Doing so allows us to estimate city-specific fixed effects, which represent the composite utility of local attributes. Second, we regress these estimated fixed effects on local amenities, using IV to correct for likely endogeneity.

We start by assuming the following utility function for individual i living in location j and consuming quantities C_i and H_i of the numeraire good and housing, respectively:

$$U_{ij} = C_i^{\beta_C} H_i^{\beta_H} X_j^{\beta_X} e^{M_{ij} + \xi_j + \eta_{ij}}. \quad (3)$$

As before, X_j denotes the local amenity of interest (here, air quality). Unobservable attributes of location j are captured in ξ_j ; η_{ij} represents an individual-specific idiosyncratic component of utility that is assumed independent of mobility costs and city characteristics. M_{ij} , meanwhile, measures the long-run (dis)utility to person i of migrating from their home state to location j . This formulation captures mobility constraints, broadly defined. An ideal specification of the model would include an explicit time dimension, tracking individual households from location to location. Data limitations prevent such an approach, however. To allow for a rich characterization of the areas where individuals might choose to live (that is, a broad sample of metropolitan areas), we rely on US Census microdata. In that data, however, only two locations are reliably observed for each individual: birth state and current residential location. As a result, our model is only partially rather than fully dynamic, taking into account where an individual was born and where they end up, but not where they might have lived along the way. Put another way, our model expands the traditional hedonic approach to take initial conditions into account.² Moreover, we focus on household heads 35 years old or younger, which further mitigates the issue. Nonetheless, our model is unable to address issues such as the riskiness of housing as an asset, and how that riskiness might affect relocation decisions.³

Individuals maximize their utility subject to the budget constraint in Eq. (1). Incorporating that budget constraint into the utility function, differentiating with respect to H_i , and rearranging yields

$$H_{ij}^* = \frac{\beta_H}{\beta_H + \beta_C} \frac{I_{ij}}{\rho_j}. \quad (4)$$

Eq. (4) states that housing expenditure accounts for a constant fraction of income, given by $\beta_H/(\beta_H + \beta_C)$ (recall that ρ_j is the price of housing services in location j). For the sake of exposition, we assume that ρ_j is known; in the empirical analysis we will estimate it from the data, as described in Section 4 below.

² For a truly dynamic model with forward-looking agents, one needs to observe individuals on multiple occasions. Ref. [5] uses restricted access NLSY geo-coded panel data to allow agents to be forward looking with respect to amenities, but is restricted by the data in the scope of the geographic choice set that agents can consider.

³ For example, households may be out of equilibrium while waiting for the housing market to make an adjustment. Their observed locations would not then accurately reflect their preferences. See [9] for a discussion of optimal behavior with respect to risky, illiquid assets like housing.

Substituting for H^* in (3) and using the budget constraint yields the indirect utility function:

$$V_{ij} = I_{ij}^{\beta_I} e^{M_{ij} - \beta_H \ln \rho_j + \beta_X \ln X_j + \zeta_j + \eta_{ij}}, \quad (5)$$

where $\beta_I \equiv \beta_C + \beta_H$. MWTP for the amenity X_j equals the marginal rate of substitution between X_j and income—i.e., for individual i , $MWTP_i = (\beta_X / \beta_I)(I_{ij} / X_j)$. Note that while the coefficient on the amenity, β_X , is constant across individuals, MWTP varies with income.

The analysis so far assumes that income I_{ij} is known for every individual in every region. In practice, of course, we must estimate what income would have been in regions not chosen. This procedure is described in Section 4. We thus decompose income into a predicted mean and an idiosyncratic error term—i.e., $I_{ij} = \hat{I}_{ij} + \varepsilon_{ij}^I$. Substituting this into Eq. (5) and taking logs yields

$$\ln V_{ij} = \beta_I \ln \hat{I}_{ij} + M_{ij} + \theta_j + v_{ij}, \quad (6)$$

where

$$\theta_j \equiv -\beta_H \ln \rho_j + \beta_X \ln X_j + \zeta_j, \quad (7)$$

and $v_{ij} \equiv \beta_I \varepsilon_{ij}^I + \eta_{ij}$; θ_j comprises all of the utility-relevant attributes of location j that are constant across individuals. Meanwhile, v_{ij} is an error term that summarizes individual i 's idiosyncratic preferences for location j .

Individuals choose their location to maximize their utility. We assume that the idiosyncratic city preferences (v_{ij}) are independently and identically distributed type I extreme value. This implies that the share of the population choosing to live in city j is given by a logit specification. Hence the probability that individual i settles in location j can be written as

$$P(\ln V_{ij} \geq \ln V_{i,l} \forall l \neq j) = \frac{e^{\beta_I \ln \hat{I}_{ij} + M_{ij} + \theta_j}}{\sum_{q=1}^J e^{\beta_I \ln \hat{I}_{i,q} + M_{i,q} + \theta_q}}. \quad (8)$$

We estimate Eq. (8) by maximum likelihood.

We recover the vector $\{\theta\}$ as parameters in the logit estimation. These city-specific fixed effects represent the indirect utility (somewhat loosely, the “quality of life”) from residing in each city, independent of mobility costs or income. In the second stage of estimation, we regress the estimated $\{\theta\}$ on local air pollution concentrations and other local amenities using Eq. (7).

2.3. Relationship between the two approaches

The discrete-choice model just outlined is closely related to the Roback model. In the latter setting, all individuals are identical (i.e., $\eta_{ij} \equiv 0$) and indifferent among locations, hence V is constant. Taking the total derivative of Eq. (5), setting it equal to zero, and treating ζ_j like another element of X_j with a coefficient equal to 1 yields

$$\frac{1}{V} \frac{dV}{dX} = \frac{\beta_I}{I} \frac{dI}{dX} + \frac{dM}{dX} + \frac{\beta_X}{X} - \beta_H \frac{d \ln \rho}{dX} = 0. \quad (9)$$

After rearranging terms and doing a bit of algebra, we have

$$\Rightarrow p^* = \frac{\beta_X}{\beta_I} \frac{I}{X} = H^* \frac{d\rho}{dX} - \frac{dI}{dX} - \frac{V_M}{V_I} \frac{dM}{dX}, \quad (10)$$

which is identical to Eq. (2). Nonetheless, the identifying assumptions of the two models are very different. The Roback model uses individuals' indifference among locations to derive the result in Eq. (10). Since, in that model, individuals equate their marginal rate of substitution with the implicit amenity price, estimating the elements of (10) amounts to inferring the MWTP.

In contrast, our discrete-choice model relies on location decisions to reveal preferences about local amenities. In our model, individuals sort among locations on the basis of idiosyncratic tastes, and thus have strict preferences over location. If we are willing to assume that a city's appeal is a weighted sum of the city's characteristics, and that the weights are constant among individuals, then we can identify the underlying MWTP directly from an equation such as (7). These additional assumptions represent the cost of our approach. The benefit is that it readily allows us to incorporate mobility costs. As we showed above, the presence of mobility costs complicates inference in the traditional hedonic model. In the empirical analysis that follows, we confirm that allowing for mobility costs makes a large difference in the estimated value of clean air.

Our discrete-choice model also highlights the question of how the size of a city should be used in inferring the value of local amenities. City size plays only an indirect role (i.e., through equilibrium housing prices and incomes) in the conventional wage-hedonic model. In contrast, our approach – by relying on residential location to reveal preferences – infers higher utility for places chosen by a larger share of individuals. All else equal, bigger cities must have larger estimated values of θ_j . If big cities are big *because* of the observable amenities they offer, then the larger estimated city-fixed effects convey useful information about how people value local attributes. On the other hand, city size might enter into individuals' utility directly (e.g., positively through agglomeration effects or negatively via congestion costs). If city size

is also correlated with local amenities (e.g., larger cities have more manufacturing facilities and thus poorer air quality), then omitting it will introduce bias. Accordingly, in our empirical analysis we report results from specifications with and without population included as a covariate.

2.4. Identification

Two final econometric issues must be addressed in estimating the second stage of our model, given by Eq. (7). First, the price of housing services, ρ_j , varies with observable characteristics of city j , and is likely correlated with unobserved local characteristics in ξ_j . We solve this problem by moving $\beta_H \ln \rho_j$ to the left-hand side of the regression equation. From Eq. (4), we have $\beta_H = \beta_i(\rho_j H_i^* / I_{i,j})$; the parameter β_i is estimated in the first stage of our procedure, and we set $\rho_j H_i^* / I_{i,j}$ (the share of housing expenditures in income) equal to its median value in our sample, which is 0.2.⁴ The new dependent variable, $\theta_j + \beta_H \ln \rho_j$, can be thought of (again somewhat loosely) as the “housing-price-adjusted quality of life,” or alternatively the net value of living in location j after accounting for housing prices.

Second, amenity levels are likely correlated with local unobservable attributes. In our case of air quality, local economic activity is likely to be positively correlated with local air pollution as well as local rents and wages. As a consequence, naïve estimation of Eq. (7) by ordinary least squares (OLS) is likely to yield biased parameter estimates. To address this potential source of bias, previous research has attempted to isolate a component of air pollution that is orthogonal to economic activity. In a recent paper, for example, [6] use discontinuities implicit in the Clean Air Act to isolate a source of pseudo-random variation in regulatory intensity across similar locations. Following [6], we combine two strategies to deal with this potential correlation. First, we estimate the modified version of Eq. (7) in first differences, using panel data from 1990 and 2000:

$$\Delta \theta_j + \beta_H \Delta \ln \rho_j = \beta_X \Delta X_j + \zeta_j, \quad (11)$$

where for example, $\Delta \theta_j = \theta_{2000} - \theta_{1990}$; and ζ_j is the time-varying component of the unobservable ξ_j . Note that we have moved $\beta_H \Delta \ln \rho_j$ to the left-hand side of the regression equation. Taking first differences eliminates any bias due to correlation between persistent air pollution and permanent unobserved city characteristics—for example, a concentration of highly polluting manufacturing industries, or perennial traffic congestion. However, one might still worry about potential correlation between ζ_j and ΔX_j .⁵ Hence we also need to find an instrument for air pollution.

We develop a novel instrument that exploits the geography of particulate matter formation and transmission. Pollution travels long distances—particulates emitted from Midwestern power plants, for example, contribute substantially to air pollution in the Northeast and Mid-Atlantic. At the same time, such emissions are likely to be uncorrelated with housing prices or local economic activity in those regions. Drawing on this intuition, we instrument for changes in local air pollution using changes in particulate matter originating from distant sources. In particular, for the years 1990 and 1999, we compute the particulate matter in location j that is attributable to all sources located at least 80 km from that location, and use the difference between the two measures as our instrument for the change in air pollution. We describe the choice of 80 km and the construction of the instrument in detail in the on-line appendix. The key step is the use of a county-to-county source–receptor matrix developed for the US EPA. This matrix relates emissions from nearly 6000 sources throughout the US to pollution concentrations in the 3080 receptor counties. By excluding sources within a chosen radius, we can construct a measure of the pollution concentration for a given city that is attributable to distant sources.

3. Data

3.1. Primary data sources

The data used for this analysis come from several sources, all publicly available. For the discrete-choice model of residential location decisions, as well as the regressions used to estimate individual income and the price of housing services at the MSA level, we draw on the 1% and 5% microdata samples of the 1990 and 2000 US Population Censuses, respectively. The census data (available at www.ipums.org) describe attributes of the household head along with the household’s composition. The data set we use for our analysis consists of random samples of 10,000 household heads in each year who are under the age of 35 and reside in one of 242 MSAs. We treat the household head as the decision maker, and focus on his/her attributes, along with those of the dwelling in which the household resides. Migration variables are calculated from data describing the household head’s state of birth and the location of each MSA. We exclude household heads over 35 years old to ensure that location decisions are driven by current local attributes. The 242 MSAs that comprise our choice set include the larger US cities, and contain approximately 86% of the total US metropolitan population in both

⁴ The estimate of 0.2 corresponds to the share of income spent on housing in our sample of individuals in the microcensus data, using a 30-year fixed mortgage rate of 9%, which is the average of the values in 1990 and 2000. In our empirical analysis, we show that our results are robust to other choices of this parameter.

⁵ Suppose, for example, that ζ_j includes the effects of an economic recession in location j . If reduced economic activity is correlated with reductions in PM pollution from reduced economic activity, the estimate of β_X may be biased upward.

Table 2
Summary statistics.

| Variable name | Description | 1990 | | 2000 | | Change | | Source |
|----------------|---|------------|-----------|--------|-----------|--------|-----------|--------|
| | | Mean | Std. dev. | Mean | Std. dev. | Mean | Std. dev. | |
| Y | Per capita income (\$000s) | 14.022 | 2.643 | 16.244 | 3.620 | 2.222 | 1.386 | (2) |
| $\ln \rho$ | $\ln(\text{Price of housing services})$ | a 3.722 | 0.430 | 4.480 | 0.337 | 0.758 | 0.269 | (1) |
| $\bar{\theta}$ | MSA-level fixed effect | a -0.006 | 1.291 | -0.004 | 1.307 | 0.002 | 0.230 | (1) |
| PM | PM10 concentration ($\mu\text{g}/\text{m}^3$) | 42.21 | 21.15 | 33.87 | 15.11 | -8.35 | 10.08 | (1) |
| $Employment$ | Fraction of population employed | 0.565 | 0.086 | 0.591 | 0.091 | 0.026 | 0.029 | (2) |
| $Manuf. est.$ | Number of manufacturing establishments | b 1078 | 2121 | 1060 | 1878 | -17 | 369 | (3) |
| $Crime$ | Crime rate per capita | c 0.350 | 0.260 | 0.289 | 0.217 | -0.061 | 0.173 | (3) |
| $Prop. tax$ | Fraction of local tax revenue from property taxes | d 0.754 | 0.166 | 0.741 | 0.162 | -0.012 | 0.059 | (3) |
| $Govt. exp.$ | Local government expenditure per capita (\$000s) | d 1.444 | 0.350 | 2.398 | 0.593 | 0.954 | 0.408 | (3) |
| $White$ | Fraction of population that is white | 0.836 | 0.103 | 0.790 | 0.113 | -0.045 | 0.029 | (3) |
| $Health$ | Health ranking | 153.20 | 91.05 | 148.20 | 89.69 | -5.00 | 43.12 | (4) |
| $Arts$ | Arts ranking | 149.71 | 89.32 | 146.35 | 89.84 | -3.37 | 52.77 | (4) |
| $Transport$ | Transportation ranking | 146.31 | 88.42 | 141.43 | 88.44 | -4.88 | 69.75 | (4) |
| $Populatoin$ | Population (millions) | 0.715 | 1.123 | 0.810 | 1.237 | 0.097 | 0.166 | (3) |

Notes:

a : The price of housing services is shown in logs to facilitate comparison with the MSA-level fixed effects.

b : Manufacturing establishments data are for years 1987 and 1997.

c : Crime rate is FBI crime rate in per capita terms for years 1990 and 1999.

d : Property tax and government expenditure data are for 1986–1987 and 1996–1997.

Sources are (1) estimates from current study, as described in text; (2) Regional Economic Information System (REIS); (3) *County and City Data Books*; and (4) *Places Rated Almanac*. All monetary values are expressed in constant 1982–1984 dollars.

1990 and 2000. The key census variables used in the analysis are described in Table A1, which is contained in an appendix that is available at JEEM's on-line archive of supplementary materials. This can be accessed at <http://www.aere.org/journal/index.html>.

To estimate MWTP for air quality, we require data on pollution, local economic activity, and a range of local amenities. We describe the construction of our air pollution measure and instruments in detail below and in the on-line appendix. Information on income, population, and employment comes from the Regional Economic Information System database maintained by the Bureau of Economic Analysis. Data on other local amenities are taken from various editions of the *County and City Data Book* and the *Places Rated Almanac* [15,16].⁶ Table 2 presents summary statistics and a full description of the variables used in the analysis.

3.2. Air quality measures

Our measure of air pollution is the ambient concentration of particulate matter.⁷ Particulate matter refers to airborne small particles, fine solids, and aerosols that form as a result of activities as diverse as the fossil fuel combustion, mining, agriculture, construction and demolition, and driving on unpaved roads. While most of the particles resulting from these processes are relatively large in size (i.e., approximately 1/7th the diameter of a human hair), smaller particles result from chemical processes that occur when sulfur dioxide, nitrogen oxides, and volatile organics react with other compounds in the atmosphere. The result is an array of pollutants, collectively known as “PM10” (because they are all smaller than 10 μm in size), which carry with them serious health consequences (see on-line appendix for details).

We estimate ambient pollution concentrations for each MSA in 1990 and 1999 using data on emissions of particulates and sulfur dioxide (a precursor to PM10).⁸ The data are taken from the National Emissions Inventory maintained by the EPA. To translate emissions into concentrations of particulate matter, we use the PM10 module of the Source–Receptor Matrix Model.⁹ [11] This procedure is described in the on-line appendix. A related procedure, also described in the on-line appendix, is used to construct our instruments.

⁶ Data from the REIS and CCDB are at the county level. We aggregate up to the metro-area level using the same MSA definitions as we use in the pollution data (based on MSA designations in 1990). Doing so ensures that our definitions of MSAs remain constant in both years, even as the official Census designations changed.

⁷ In an ideal world, we would estimate our model using additional measures of ambient air pollution, such as sulfur dioxide (SO_2) or ground-level ozone (O_3). However, we are unaware of any fine-grained S–R matrix (which we require for our instrumenting strategy) for other pollutants comparable to the PM10 S–R matrix we use here. A consolation is that PM10 is far and away from the most important air pollutant in terms of human health effects (see the discussion in the on-line appendix) and will be highly correlated with other pollutants like SO_2 .

⁸ We use data for 1999 rather than 2000 because the National Emissions Inventory is collected at 3-year intervals.

⁹ We thank Wayne Gray and his co-authors for generously sharing the S–R matrix with us. The discussion of the matrix is based in part on the discussion in [17]. The report by [1] also gives a detailed exposition of the S–R matrix we use here.

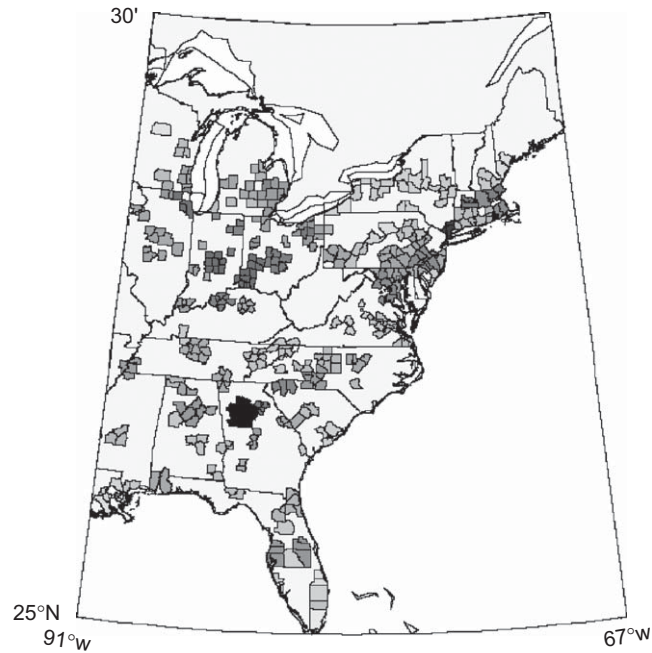


Fig. 1. Computed PM10 concentrations in the eastern United States.




Fig. 2. Computed PM10 concentrations in the western United States.

Figs. 1 and 2 illustrate our pollution data, depicting computed PM10 concentrations in 1999 for each of the 242 MSAs in our data. Darker shadings correspond to greater concentrations of particulate matter. The western and eastern United States are depicted separately, but the same shading gradient is used. Note that ambient concentrations of particulates generally increase from west to east, mirroring the underlying weather patterns. Thus the cities of the West Coast have relatively low levels of particulates on the whole, while the nation's highest concentrations occur in Atlanta and New York.

4. Econometric specification

Several steps are needed to implement the residential sorting model outlined in Section 2. First, we must estimate housing prices and incomes in each location. Next, we must choose a representation of mobility costs. We can then use a

logit model of location choice to estimate the city-specific fixed effects. Finally, we regress those fixed effects on local attributes. We discuss the details of each step in turn. Throughout this analysis, we use i to index households, j to index locations (MSAs), and t to index the year (1990 or 2000). We will often pool data from both years. Note that while the set of metropolitan areas is the same in each year, the set of households is not.

One approach to quantifying housing prices would be to take an aggregate measure – for example, the median value of a home in each MSA. However, such an approach raises potentially serious problems of aggregation bias. In particular, home values might rise because of unobserved changes in the quality of the housing stock, rather than changes in local amenities. If these changes in housing supply are correlated with local amenities, an endogeneity problem arises.¹⁰

We employ a different approach that takes explicit account of the characteristics of individual homes. Let $P_{i,j,t}$ denote the value of the home owned by household i in location j appearing in year t , which we define as the value of the house (for owner-occupied housing) or monthly rent (for rental units). We model $P_{i,j,t}$ as a function of the characteristics of the dwelling, given by a vector $\mathbf{h}_{i,t}$, and a scaling parameter $\rho_{j,t}$ specific to city j and year t . $\Omega_{i,t}$ is a dummy variable that equals 1 if household i owns its home and 0 otherwise; thus $\lambda_{j,t}$ measures the premium on owned housing. Taking logs:

$$\ln P_{i,j,t} = \ln \rho_{j,t} + \lambda_{j,t} \Omega_{i,t} + \mathbf{h}'_{i,t} \boldsymbol{\phi}_t + e^H_{i,j,t}. \tag{12}$$

Along with housing characteristics, the parameters $\boldsymbol{\phi}_t$ yield an index of “housing services” each period, defined as $H_{i,t} = \exp(\mathbf{h}'_{i,t} \boldsymbol{\phi}_t)$. Hence the parameter $\rho_{j,t}$ measures the effective “price of housing services” in a particular location and a particular year. Because we control for the bundle of housing services, these prices provide a consistent measure of the true price of housing across metropolitan areas with different housing stocks. We can readily estimate these prices as the MSA and time-specific intercepts in a regression of Eq. (12), using the census microdata described in Section 3.

Next, consider income. We do not observe the income that a given individual would earn in every location, but only what he earns in his chosen city. In the microdata used for estimation, however, household heads with similar characteristics are scattered among locations. Hence we can compute the income each individual would earn in every location by estimating a series of location-specific regressions of incomes on a set of individual attributes, controlling for non-random sorting as in [7]. See the on-line appendix for details.

Next, consider mobility costs. Table 1 suggests that households tend to settle close to where the household head was born. We capture this feature of the data with a flexible migration cost matrix:

$$M_{i,j,t} \equiv f_M(\mathbf{d}_{i,j,t}; \boldsymbol{\mu}) = \mu_S d^S_{i,j,t} + \mu_{R1} d^{R1}_{i,j,t} + \mu_{R2} d^{R2}_{i,j,t}, \tag{13}$$

where $d^S_{i,j,t} = 1$ if location j is outside individual i 's birth state ($= 0$ otherwise), $d^{R1}_{i,j,t} = 1$ if location j is outside individual i 's birth region as defined in Table 1 ($= 0$ otherwise), and $d^{R2}_{i,j,t} = 1$ if location j is outside individual i 's macro-region ($= 0$ otherwise).¹¹ We normalize migration costs to zero if the household head does not leave his birth state.

We now turn to the estimation of the parameter vector $\{\mu_S, \mu_{R1}, \mu_{R2}, \beta_1, \boldsymbol{\theta}\}$. On the assumption that preferences are stable over time, we can estimate a single set of mobility parameters μ and marginal utility of income β_1 for both years 1990 and 2000. We do this by pooling the data over the two years, and calculating a single likelihood function:

$$L(\mu_S, \mu_{R1}, \mu_{R2}, \beta_1, \boldsymbol{\theta}) = \prod_t \prod_i \prod_{j=1}^J \left[\frac{e^{\beta_1 \ln \hat{i}_{i,j,t} + \mu_S d^S_{i,j,t} + \mu_{R1} d^{R1}_{i,j,t} + \mu_{R2} d^{R2}_{i,j,t} + \theta_{j,t}}}{\sum_{k=1}^J e^{\beta_1 \ln \hat{i}_{i,k,t} + \mu_S d^S_{i,k,t} + \mu_{R1} d^{R1}_{i,k,t} + \mu_{R2} d^{R2}_{i,k,t} + \theta_{k,t}}} \right]^{\chi_{i,j,t}}, \tag{14}$$

where $\chi_{i,j,t}$ is an indicator function that equals one if household i observed in year t chooses location j , and zero otherwise.^{12,13}

Recall that $\{\theta_{j,t}\}$ represent composite city-level attributes. Let $PM_{j,t}$ denote the air pollution (PM10) concentration in location j and period t , computed as described in the on-line appendix. Note that higher values of $PM_{j,t}$ correspond to worse air quality, so that $\beta_{PM} < 0$ if individuals are willing to pay for better air quality. Let $\mathbf{Z}_{j,t}$ denote a vector of other observable city attributes. The equation to be estimated in the second stage is thus (updating Eq. (11))

$$\Delta \theta_j + 0.2 \Delta \ln \rho_j = \beta_{PM} \Delta \ln PM_j + \Delta \mathbf{Z}'_j \boldsymbol{\beta}_Z + \zeta_j. \tag{15}$$

We estimate Eq. (15) by IV, using particulate matter concentrations based on sources outside 80 km ($\Delta \ln PM_{Rj}^{80}$) interacting with regional dummies, as our instrument for $\Delta \ln PM_j$. The covariates in $\mathbf{Z}_{j,t}$ include a range of local characteristics of metropolitan areas, including local economic activity, crime, local government tax and expenditure data, and rankings of MSAs in various categories of quality of life such as health care provision, arts, and transportation infrastructure.

¹⁰ Ref. [6] uses median home prices at the county level as its dependent variables in hedonic estimation, controlling for the potential bias by including a range of county-level characteristics of the housing market. As it argues, its instrumental variables strategy should also help eliminate the bias.

¹¹ There are four macro-regions defined by the US Census Bureau: (1) Northeast, (2) Midwest, (3) South, and (4) West.

¹² In practice, when the choice set is large (as it is in our application), estimating the full vector $\boldsymbol{\theta}$ by maximum likelihood can be computationally prohibitive. Ref. [4] provides a computational algorithm whereby these values are imputed indirectly.

¹³ Note that there is an arbitrary normalization of one of the $\theta_{j,t}$ values: raising the utility of all locations by a constant amount leaves location decisions unchanged. We set $\theta_{j,t}$ equal to zero for the Houston, TX, MSA.

Table 3
Results from conventional wage-hedonic regressions.

| Dependent variable | OLS | | IV | |
|--------------------|----------------------|----------------------|----------------------|----------------------|
| | (1) | (2) | (3) | (4) |
| $\Delta \ln \rho$ | -0.232** (0.097) | -0.292*** (0.098) | -0.497*** (0.179) | -0.634*** (0.185) |
| $\Delta \ln Y$ | -0.073*** (0.022) | -0.074*** (0.023) | -0.035 (0.041) | -0.006 (0.043) |
| MSA covariates | No | Yes | No | Yes |
| Regional dummies | Yes | Yes | Yes | Yes |

Notes: This table presents results from conventional wage-hedonic regressions. The cells contain the coefficients on $\Delta \ln(\text{PM10})$ pertaining to housing services (ρ) and income (Y) with respect to increases in air pollution. Columns (1) and (2) present results from OLS regressions; columns (3) and (4) present results using estimated PM10 from sources farther than 80 km as an instrument. Standard errors are in parentheses; * denotes significance at 10%; ** at 5%; *** at 1%.

5. Estimation results

5.1. Housing-price and income regressions

Results from the housing-price regressions described in Eq. (12) are reported in Table A2 of the on-line appendix for each year. Results are as expected. Bigger, newer houses yield more housing services, as do houses on larger plots and with complete kitchen and plumbing facilities. An inspection of the most and least expensive cities in the US in terms of the price of housing services corresponds to conventional wisdom. The average price of housing services more than doubles between 1990 and 2000, while the premium on owned housing rises by 14%. All estimates are statistically significant at the usual levels.

Table A3, also in the on-line appendix, summarizes the results from the MSA-specific income regressions. Men earn more than women, whites earn more than minorities, and income increases with education. Income falls significantly for those over age 60, reflecting retirement patterns. The premiums for white and male, along with the age penalty, all diminish between 1990 and 2000. In the case of age penalty, this fall may reflect growing participation in the labor market after age 60. Over the same time period, the premium for college education rises, while that for a high school diploma falls.

5.2. Estimates from the conventional model

As a benchmark for comparison with our residential sorting model, we estimate a conventional hedonic model without mobility costs. We estimate the model with and without instruments for air pollution, as a preliminary assessment of the severity of the bias and the success of our instrumenting strategy.

Recall that in the conventional model with costless migration, the implicit price of local amenities – and hence MWTP – can be estimated as the sum of the housing-price and income gradients with respect to a given amenity. Accordingly, we regress the log of per-capita income in MSA- j (denoted Y_j) and the price of housing services ρ_j (the MSA-specific intercept from the housing-price regressions) on particulate matter concentrations PM_j and the matrix of regional dummies \mathbf{R} . We estimate these equations in first differences:

$$\ln \Delta Y_j = \gamma_{PM,Y} \Delta \ln PM_j + \Delta \mathbf{Z}_j \boldsymbol{\beta}_Z + \gamma_{R,Y} R_j + u_j^Y, \quad (16)$$

$$\ln \Delta \rho_j = \gamma_{PM,\rho} \Delta \ln PM_j + \Delta \mathbf{Z}_j \boldsymbol{\beta}_Z + \gamma_{R,\rho} R_j + u_j^\rho. \quad (17)$$

Table 3 reports OLS and instrumental variable estimates of the coefficients on PM10 concentrations, i.e., $\gamma_{PM,Y}$ and $\gamma_{PM,\rho}$.¹⁴ Both OLS estimates are negative and significantly different from zero. Taken at face value, these coefficients imply that both

¹⁴ For housing-price regression, the covariates in the specifications reported in columns (2) and (4) are the same as in the main results from the residential sorting model (refer to Table 5). For the income regression, the employment rate and number of manufacturing establishments are omitted because they are simultaneously determined with wages and salaries.

Table 4

Results from first-stage discrete choice model of residential location decision.

| Variable | Parameter | Coefficient | t-Statistic |
|----------------------------|------------|-------------|-------------|
| Migration cost | | | |
| State | μ_S | -2.900 | -22.0 |
| Region | μ_{R1} | -0.855 | -11.5 |
| Macro-region | μ_{R2} | -0.591 | -12.5 |
| Marginal utility of income | β_I | 0.673 | 48.4 |

Table 5

Results from second-stage regressions.

| Dependent variable | OLS | | IV | | |
|----------------------------------|----------------|-------------------|------------------|-------------------|-------------------|
| | (1) | (2) | (3) | (4) | (5) |
| $\Delta\theta+0.25\Delta\ln\rho$ | | | | | |
| $\Delta\ln(PM)$ | -0.086 (0.060) | -0.107** (0.054) | -0.255** (0.110) | -0.286*** (0.109) | -0.230** (0.101) |
| $\Delta\ln(Crime)$ | | 0.010 (0.067) | | 0.008 (0.073) | 0.024 (0.068) |
| $\Delta\ln(Prop. tax)$ | | 0.359* (0.186) | | 0.396* (0.203) | 0.346* (0.188) |
| $\Delta\ln(Govt. exp.)$ | | 0.112*** (0.039) | | 0.131*** (0.042) | 0.114*** (0.039) |
| $\Delta\ln(White)$ | | -0.064 (0.389) | | -0.132 (0.424) | -0.034 (0.394) |
| $\Delta\ln(Health)$ | | -0.001*** (0.000) | | -0.001*** (0.000) | -0.001*** (0.000) |
| $\Delta\ln(Arts)$ | | 0.000 (0.000) | | 0.000 (0.000) | 0.000 (0.000) |
| $\Delta\ln(Transport)$ | | 0.000 (0.000) | | 0.000 (0.000) | 0.000 (0.000) |
| $\Delta\ln(Employment)$ | | -0.367 (0.391) | | -0.612 (0.424) | -0.319 (0.397) |
| $\Delta\ln(Manuf. est.)$ | | 0.023 (0.087) | | 0.271*** (0.081) | 0.020 (0.088) |
| $\Delta\ln(Population)$ | | 0.820*** (0.146) | | | 0.823*** (0.148) |
| Constant | -0.020 (0.050) | -0.058 (0.053) | -0.087 (0.063) | -0.049 (0.065) | -0.101 (0.061) |
| Regional dummies | Yes | Yes | Yes | Yes | Yes |
| R ² | 0.08 | 0.32 | 0.05 | 0.19 | 0.31 |
| Observations | 242 | 242 | 242 | 242 | 242 |

Notes: Standard errors in parentheses.

* Significance at 10%.

** Significance at 5%.

*** Significance at 1%.

housing prices and wages rise when air quality improves. The former effect is consistent with expectations, but not the latter.

The effects of instrumenting for air pollution suggest that the OLS estimates are indeed biased. The estimated elasticity of housing prices with respect to air pollution more than doubles in magnitude, going from -0.30 in the OLS regression to -0.63 in the IV estimates. Meanwhile, the effect of air pollution on per-capita income vanishes. We ignore the results from the income equation in computing MWTP.¹⁵ This is a conservative approach. Since the estimates are negative, excluding them inflates the estimates of MWTP from the conventional model. This closes the gap between those estimates and those from our discrete-choice model below, leading us to understate the importance of migration costs.

5.3. Estimates from the residential sorting model

Table 4 reports parameter estimates from the first-stage residential choice Eq. (14). Estimates are highly statistically significant and have the expected signs. There is a significant utility cost associated with leaving one's birth state. Costs continue to rise with leaving one's birth region and macro-region, but at a declining rate. The estimate of the marginal utility of income (β_I) is 0.673. The results of primary interest from the first stage, of course, are the MSA-level fixed effects. While too numerous to summarize, these can be illustrated with some examples. Controlling for population, the three least attractive metropolitan areas in the year 2000 (those with the most negative values of θ_j) were New Bedford, MA; Danbury, CT; and Detroit, MI. Cities near the median included Memphis, TN (#116 out of 242) and Hartford, CT (#125). Portland, OR; and Providence, RI; ranked among the top five.¹⁶

These estimated MSA-level fixed effects are used as the dependent variables in the second-stage estimation of Eq. (15). Table 5 reports results for a range of specifications. Columns (1) and (2) present OLS estimates; columns (3)–(5) report

¹⁵ Ref. [6] reports a similar finding, and likewise ignores the income estimates in computing MWTP.

¹⁶ In the raw rankings, city size makes a big difference, as we discussed in Section 2.3. Without controlling for population, the cities with the highest estimates of $\theta_{j,t}$ are Los Angeles, Chicago, and New York. Of course, controlling for population has a much smaller effect on the change from 1990 to 2000.

results from IV estimation. To account for the potential role played by city size, we include the logarithm of population as a covariate in the specifications reported in columns (2) and (5).

The estimated coefficients on $\Delta \ln(PM)$ are presented in the first row of the table. Dividing these coefficients by the marginal utility of income (0.673) yields the elasticity of willingness to pay with respect to air pollution concentrations. As in the housing-price regressions (Table 3), OLS yields statistically significant estimates with the expected sign. Once again, a comparison with the IV results reveals strong evidence of endogeneity bias. When we instrument for air pollution, the estimated elasticity nearly triples in magnitude—the OLS estimates are -0.13 to -0.16 , while the IV estimates range from -0.34 to -0.42 .

Note that these willingness-to-pay elasticities are not directly comparable to the elasticities of housing prices reported in Table 3. The estimates from the conventional model represent the percent change in housing expenditures associated with a 1% change in air pollution. In contrast, the elasticities estimated by the residential sorting model incorporate not only changes in housing prices, but also foregone income and the disutility from moving. Thus the relative magnitudes of the raw parameter estimates in the conventional model are misleading; in dollar terms, as we shall see below, the results from the residential sorting model are more than three times as large as the results from the conventional approach. Hence including mobility costs matters greatly for our estimates.

For some perspective on what these estimates imply, consider a concrete example. In 1990, the computed PM10 concentration in the New Haven–Meriden MSA was $62.2 \mu\text{g}/\text{m}^3$. In the same year, the computed PM10 concentration in the Raleigh–Durham–Chapel Hill MSA was $44.0 \mu\text{g}/\text{m}^3$ —roughly 30% lower than in New Haven, or almost exactly one standard deviation away (the standard deviation of PM10 in the sample is $18.8 \mu\text{g}/\text{m}^3$). The estimated elasticity in the full specification is -0.34 (i.e., dividing the appropriate number in Table 5, column 5 by the marginal utility of income), implying that the increase in air quality on moving from New Haven to Durham would correspond to an increase in willingness to pay of 10%. Per-capita income in 1990 in New Haven was \$23,558 (in current dollars). Hence the air quality benefits of moving from New Haven to Durham were worth roughly \$2360 in foregone consumption.

The estimated coefficients on other local amenities included in the regressions (the covariates in Z) vary in significance, but most have the expected signs in the IV estimates.¹⁷ Importantly, the inclusion of metropolitan area characteristics in general has only a small effect on the estimated coefficient on $\Delta \ln PM$. This robustness provides additional support for our IV strategy. Finally, the coefficient on population is highly significant, in line with expectations. Controlling for population reduces the magnitude of the estimated impact of air quality by just under 20%. Thus the specifications with and without population define the range of MWTP.

The coefficients on these covariates allow the computation of elasticities with respect to local amenities other than air quality. For example, the median value of the health care ranking (across both years) is 145.5. Thus the estimated coefficient of -0.001 on $\Delta(\text{Health})$ from the full specification implies that a 1% increase in a city's health care ranking translated into an approximate 0.1% increase in its attractiveness, as measured by willingness to pay. Similarly, the elasticity of willingness to pay with respect to government spending (at the sample median of 1.34, in thousands of dollars) is 0.23. By comparison, the elasticity with respect to air quality is larger but of the same general magnitude. Alternatively, we can estimate the percentage increase in willingness to pay that resulted from the median change in local amenities in the data. For government expenditures, the median change was 0.22, corresponding to a 4% increase in willingness to pay. The median change in PM concentration was a reduction of $5.7 \mu\text{g}/\text{m}^3$, which translates into a 5% increase in willingness to pay. Thus the median improvement in air quality was comparable, in quality-of-life terms, to the median increase in per-capita government expenditure.

Table A5 in the on-line appendix presents estimated PM10 utility parameters under alternative assumptions about (i) the share of expenditures devoted to housing, (ii) the exclusion distance used in constructing the instrument for PM10, and (iii) functional form. Based on this evidence, we conclude that our base specification is a reasonable one, and that our conclusions are robust to the choice of empirical strategy in the second stage.

5.4. Marginal willingness to pay

In our residential sorting model, β_{PM} divided by β_I measures the elasticity of willingness to pay with respect to PM10 concentrations; thus we can recover an estimate of MWTP for air quality, in dollar terms, by multiplying this elasticity by household income and dividing by the PM10 concentration. To calculate a comparable MWTP for the conventional wage-hedonic model, we must first account for the share of expenditure in housing. Since the wage-hedonic model ignores mobility costs, it assumes that individuals' entire MWTP is captured in housing prices and incomes. To translate the estimates of housing-price elasticities into MWTP, therefore, we must first multiply the coefficients from the conventional

¹⁷ Metropolitan areas with higher government expenditure per capita are significantly more appealing; the fraction of revenue raised by property taxes also has a positive effect. Areas with better health care attract more residents; note that a positive value for $\Delta(\text{Health ranking})$ corresponds to a worsening of health care (and likewise for the arts and transport variables). Culture and transportation are also valued, although the estimates are imprecise. In the specification of column (4), the size of the local economy (as measured by manufacturing establishments) is positively and significantly correlated with the appeal of an area, but the effect vanishes when population is included (the two variables are strongly correlated). Our other measure of local economic activity (i.e., employment as a fraction of the total population) turns out to be insignificant.

Table 6
Estimated marginal willingness to pay for air quality.

| Measure | Hedonics | | Residential sorting | | | |
|----------------|------------------------|------------------------|------------------------|------------------------|-------------------|-------------------------------|
| | OLS | IV | OLS | IV | | |
| | Full specification (1) | Full specification (2) | Full specification (3) | Full specification (4) | No covariates (5) | No control for population (6) |
| WTP Elasticity | 0.06 | 0.13 | 0.16 | 0.34 | 0.38 | 0.42 |
| MWTP (\$) | 25.40 | 55.20 | 69.10 | 148.70 | 164.72 | 184.89 |

Notes: Specifications (1)–(4) are full specifications. Specification (5) includes no covariates. Specification (6) includes no control for population. “Hedonics” coefficients are taken from the wage-hedonic model summarized in Table 4 (columns 2 and 4). “Residential sorting” coefficients are taken from Table 5; columns 3–6 above correspond to columns 2, 5, 3, and 4 in Table 5, respectively. Marginal willingness to pay (MWTP) is calculated by multiplying the regression coefficients by the median household income in constant 1982–1984 dollars (\$15,679) and dividing by the median PM10 concentration in the sample ($36.0 \mu\text{g}/\text{m}^3$). Figures for the wage-hedonic model exclude the estimated effects of PM10 on income, which were insignificant in the IV model. All estimates are in constant 1982–1984 dollars.

hedonic model by the share of income devoted to housing expenditures, or 0.2. Multiplying the resulting elasticity by income, and dividing by the PM10 concentration, yields estimated MWTP, just as in the residential sorting model.

Table 6 reports the results of these computations, using the median values of household income (\$15,679) and PM10 concentration (36.0) in our sample as our measures of income and air pollution, respectively. Thus the reported figures for MWTP represent the median household’s willingness to pay for a $1 \mu\text{g}/\text{m}^3$ reduction in ambient PM10 concentrations, expressed in constant 1982–1984 dollars.¹⁸ For reference, we have also included the elasticities estimated from the regressions (expressed in terms of air quality rather than air pollution).

The results provide striking evidence of the importance of accounting for endogeneity and mobility costs. When we instrument for air pollution in the full model, the estimated MWTP more than doubles, increasing from \$69 to \$149. Incorporating mobility costs matters even more. The MWTP estimated by our residential sorting model is much larger than the comparable estimate from the conventional hedonic model—MWTP increases from \$55 to \$149 in comparable specifications (compare columns (3) and (4) of Table 6). We also present MWTP for the IV specifications with other sets of covariates (columns (3) and (4) in Table 5). The estimated elasticity of 0.38 from the specification without MSA covariates implies a MWTP of \$165. When local attributes other than population are included, estimated MWTP rises to \$185.

Thus a broad statement of our results is that we find an estimated MWTP for air quality ranging from \$149 to \$185, in constant 1982–1984 dollars, for a household whose head earns the median income of \$15,679. These estimates are large relative to the previous hedonics literature; for example, [6] reports an elasticity of housing prices with respect to particulate matter concentrations of -0.20 to -0.35 – i.e., half as large as our conventional hedonic estimates, and roughly one sixth the size of the elasticities estimated by our residential sorting model.

The discrepancy is even larger in dollar values. Our model identifies MWTP in terms of foregone consumption of housing services and other goods. As a result, our estimates correspond to *annual* MWTP—equivalently, the willingness to pay for a one-unit improvement in air quality that lasts for 1 year. The comparable estimates in [6] correspond to a MWTP of \$22 for a reduction in PM10 concentrations—one seventh the size of our lower estimate.¹⁹ Moreover, the estimates in [6] were themselves much larger than in the previous literature. Part of the discrepancy between their estimates and our estimates using the conventional hedonic approach can be explained by rising willingness to pay between the 1970s and 1990s due to rising incomes; [18] reports finding such an effect in their comparative analysis of the previous literature. Moreover, our MWTP estimates likely capture the effects of other pollutants whose concentrations are correlated with PM10.

The internal comparison between our estimates from the wage-hedonic model and the residential sorting model remains striking. Incorporating mobility costs yields estimates of MWTP that are more than three times as large as estimates from a conventional model. In other words, assuming that migration is costless would result in understating willingness to pay for air quality by roughly two thirds.

6. Conclusions

This paper argues that mobility constraints hinder the use of conventional wage-hedonic techniques to estimate household MWTP for local amenities such as clean air. We develop and implement a discrete-choice model that uses data

¹⁸ Measurement in 1982–1984 dollars facilitates comparison with the numbers reported by [6,18].

¹⁹ Like the literature before them, [6] frame their results in terms of total suspended particulates (TSP), which was the preferred measure of particulate pollution prior to 1987. In order to convert our WTP estimates to results in terms of TSP, one should divide our measures by approximately 1.82.

on residential patterns, along with a flexible model of migration costs, to infer the utility of living in individual metropolitan areas across the US. We then estimate the MWTP for a reduction in air pollution, as measured by the ambient concentration of particulate matter (PM10), using the contribution of distant sources to local air pollution as an instrument.

Our results suggest that the conventional approach (ignoring mobility costs) substantially understates the true MWTP for air quality. Our estimates imply that the median household would pay \$149–\$185 for a one-unit reduction in PM10 concentrations, in constant 1982–1984 dollars. These estimates are three times as large as the corresponding estimate of MWTP from a conventional hedonic model estimated using the same data. Instrumenting for local air pollution makes a large difference in both models.

These findings highlight the potential importance of incorporating mobility constraints into hedonic models. We suspect that the larger the costs relative to the benefits at stake, the greater the consequences of ignoring mobility costs. For example, while households value clean air, few are likely to leave behind their hometowns and families purely for the sake of modest reductions in air pollution. More generally, mobility costs are more likely to constrain choices among metropolitan areas, rather than among neighborhoods within a metropolitan area.

The adverse health impacts of air pollution have prompted a wide array of legislative responses at both the state and federal levels over the last 30 years. Evaluated according to simple criteria (i.e., emissions reductions and cost-effectiveness), these policies are generally considered to have been successful. Even so, studies find that over 81 million Americans face unhealthy short-term exposure to PM, while 66 million live with chronically high exposure [2]. This is cause for concern, particularly in light of current legislative efforts that would reduce the capacity of the EPA to regulate certain pollution sources (i.e., new power plants). While most of these legislative efforts arise out of concern for the cost of compliance with EPA regulations, little is known about the size of the benefits. This complicates careful evaluation based on efficiency criteria. The present study suggests that the true value of clean air may be substantially greater than has been recognized.

Acknowledgments

We would like to thank Wayne Gray, Michael Greenstone, Brian Murray, Sharon Oster, Ray Palmquist, Daniel Phaneuf, Kerry Smith, and seminar and conference participants at Brown, Harvard–MIT, Yale, Resources for the Future, the 2004 AERE Summer Workshop, and the 2005 Camp Resources for their helpful comments and insights. All remaining errors and omissions are our own.

Appendix A. Supplementary materials

Supplementary data associated with this article can be found in the on-line version at [doi:10.1016/j.jeem.2008.08.004](https://doi.org/10.1016/j.jeem.2008.08.004).

References

- [1] Abt Associates. The particulate-related health benefits of reducing power plant emissions, Prepared for Clean Air Task Force, October 2000.
- [2] American Lung Association, State of the Air: 2004, American Lung Association, Washington, DC, 2004.
- [3] P. Bajari, L. Benkard, Demand estimation with heterogeneous consumers and unobserved product characteristics: a hedonic approach, *J. Polit. Economy* 113 (6) (2005) 1239–1276.
- [4] S. Berry, Estimating discrete choice models of product differentiation, *RAND J. Econ.* 25 (1994) 242–262.
- [5] K. Bishop, A dynamic model of location choice and hedonic valuation, Working Paper, Duke University Department of Economics, 2008.
- [6] K.Y. Chay, M. Greenstone, Does air quality matter? Evidence from the housing market, *J. Polit. Economy* 113 (2) (2005) 376–424.
- [7] G. Dahl, Mobility and the return to education: testing a Roy model with multiple markets, *Econometrica* 70 (6) (2002) 2367–2420.
- [8] I. Ekeland, J. Heckman, L. Nesheim, Identification and estimation of hedonic models, *J. Polit. Economy* 112 (1.2) (2004) S60–S109.
- [9] S.J. Grossman, G. Laroque, Asset pricing optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* 58 (1990) 25–51.
- [10] D. Harrison, D.L. Rubinfeld, Hedonic housing prices and the demand for clean air, *J. Environ. Econ. Manage.* 5 (1978) 81–102.
- [11] D.A. Latimer, Particulate matter source–receptor relationships between all point and area sources in the United States and PSD Class I area receptors, Prepared for EPA, Office of Air Quality Planning and Standards, September 1996.
- [12] J. Roback, Wages, rents, and the quality of life, *J. Polit. Economy* 90 (1982) 1257–1278.
- [13] S. Rosen, Hedonic prices and implicit markets: product differentiation in pure competition, *J. Polit. Economy* 82 (1974) 34–55.
- [14] R.G. Ridker, J.A. Henning, The determinants of residential property values with special reference to air pollution, *Rev. Econ. Statist.* 49 (1967) 246–257.
- [15] D. Savageau, B. Richard, Places Rated Almanac, MacMillan/Prentice Hall Travel, New York, NY, 1993.
- [16] D. Savageau, D. Ralph, Places Rated Almanac: Millennium Edition, Hungry Minds Inc., New York, NY, 2000.
- [17] R.J. Shadbeigian, W. Gray, C. Morgan, Benefits and costs from sulfur dioxide trading: a distributional analysis, in: R.V. Gerald, M.W. Diana (Eds.), *Acid in the Environment: Lessons Learned and Future Prospects*, Springer Science+Business Media, New York, 2007.
- [18] V.K. Smith, J. Huang, Can markets value air quality? A meta-analysis of hedonic property value models, *J. Polit. Economy* 103 (1995) 209–227.