Essays on Public and Labor Economics

Essays on Public and Labor Economics

A dissertation presented

by

Daniel McArthur Sullivan

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

May 2016

Dissertation Advisor: **Professor David Cutler** **Daniel McArthur Sullivan**

Essays on Public and Labor Economics

Abstract

Chapter 1 presents evidence that current economics research significantly underestimates the effects of air pollution, regardless of the outcome of interest. This bias exists even in quasi-experimental estimates and arises because popular methods used by economists, including geographic diff-in-diffs and monitor-based interpolations, are unable to account for sharp changes in exposure over short distances. To solve this problem, I use an atmospheric dispersion model to determine the effect of every polluting firm on every house in greater Los Angeles. I then estimate the effect of NO_x emissions on house prices using the exogenous variation in emissions caused by the California Electricity Crisis of 2000 and a cap-and-trade program in greater Los Angeles. The estimated price response is much larger than past estimates while conventional methods are unable to detect any effect.

In Chapter 2, I use these methods to explore the equity implications of the Crisis-induced pollution reduction. I also present a locational equilibrium model and derive conditions under which lower-income residents are displaced by higher-income immigrants after an arbitrary local amenity is improved. I find that rents increased significantly in improved neighborhoods, on par with house prices. Simultaneously, total population decreased, driven by a mass outmigration of low-education residents. Low home-ownership rates among low-income households suggests that emigrants were not responding to a wealth windfall but were instead made worse off by the amenity improvement.

Chapter 3 considers the principal-agent problem that arises when consumers file for bankruptcy. Lawyers advise debtors on whether to file the cheaper Chapter 7 filing or the more expensive Chapter 13 filing. Bankruptcy courts that allow lawyers to charge more for Chapter 13 see a significantly larger fraction of Chapter 13 filings. This is true controlling for a host of demographic controls at the zip code level and with state fixed effects and district policy controls. Our estimates suggest that 5.4% of cross-district variation in relative Chapter 13 rates could be eliminated by harmonizing relative fees.

Contents

List of Tables

List of Figures

Acknowledgments

For Chapters 1 and 2, I am grateful to my advisers, David Cutler, Edward Glaeser, Larry Katz, and Robert Stavins, for their feedback. I also thank John Coglianese, Timothy Layton, Jing Li, Jonathan Libgober, Amanda Pallais, Christopher Palmer, Jisung Park, Daniel Pollmann, and James Stock, as well as seminar participants at Harvard, BYU, Notre Dame, the University of Wisconsin at Madison, Resources for the Future, and the NBER. Funding from the National Institute on Aging, through Grant Number T32-AG000186 to the National Bureau of Economic Research, is gratefully acknowledged.

For Chapter 3, I thank Lars Lefgren, Val Lambson, David Sims, Joseph Price, Mike Ransom and Mark Showalter for helpful comments and suggestions and Andreas Lehnert and Michelle Miller for comments and sharing data.

Most of all, I thank my wife, Brittney Stafford Sullivan, without whom this dissertation would never have been written.

To Mom, Dad, Brittney, and Grace

Chapter 1

The True Cost of Air Pollution: Evidence from the Housing Market

1.1 Introduction

House price capitalization is routinely used to measure the social value of local amenities which lack an explicit market. But the case of air pollution presents a puzzle: The house price response to improved air quality is surprisingly weak compared to the expected health benefits (Smith and Huang 1995). Studies since Smith and Huang (1995), including quasi-experimental studies, have not resolved this discrepancy.[1] Even more puzzling, such a discrepancy exists only for pollution—it is absent for school quality (Black 1999; Cellini, Ferreira, and Rothstein 2010), crime risk (Linden and Rockoff 2008; Pope 2008), and local cancer risk (Davis 2004). And it cannot be explained by public ignorance of or indifference to the dangers of pollution (Neidell 2009; Moretti and Neidell 2011).

In this paper, I show that this puzzle is a special case of a larger problem in the economics

1. See Section 1.2. For a family of two adults and one child, 1 $\mu g/m^3$ of particulate matter ($PM_{2.5}$) costs about $1,600 in increased mortality risk alone, to say nothing of acute illness risk, decreased quality of life, or the costs borne by other family members (see footnote 3). Chay and Greenstone (2005) estimate a MWTP of $191 per 1 $\mu g/m^3$ of TSP reduction. While assuming costless moving, Bayer, Keohane, and Timmins (2009) estimate MWTP of $130 for 1 $\mu g/m^3$ of PM_{10}. Currie and Walker (2011) find no significant effect on prices after a drop in NO_x and CO near highway toll booths.

literature on air pollution: Estimates of pollution's effects are systematically biased by how pollution exposure to individuals is measured. Pollution concentrations spike sharply downwind of pollution sources, with sharp changes near the source itself, e.g., immediately downwind and upwind of a highway. This undermines the most common econometric tools used to study pollution—geographic difference-in-differences and monitor-based interpolation—because they are unable to capture this sharp geographic variation. This leads to contaminated treatment and control groups and/or non-classical measurement error. The result is biased estimates, even with an exogenous shock to pollution exposure and regardless of the outcome being studied.

In turn, these biased estimates lead to understated valuations of air quality regulations and other programs like subsidies to clean energy.

I solve this problem using tools from atmospheric chemistry and show that house prices respond dramatically to changes in air quality and that conventional methods are unable to detect this response. Specifically, I measure local exposure using AERMOD, an atmospheric dispersion model developed by the American Meteorological Society and the EPA. AERMOD uses data on meteorology (e.g., wind and temperature at multiple altitudes, pressure, surface roughness) and firms (e.g., smoke stack height, diameter, gas temperature) to determine where pollutants go after leaving a firm's smoke stack. To resolve the usual concerns about the endogeneity of local pollution exposure and housing decisions, I exploit the California Electricity Crisis of 2000 as a natural experiment. The Crisis unexpectedly and permanently lowered NO_x emissions in southern California by precipitating the near collapse of RECLAIM, a then-nascent cap-and-trade market for NO_x, which hastened and synchronized firms' adoption of abatement technology.

Using AERMOD to determine who benefited from the Crisis reveals that the demand for clean air is high while using conventional methods does not. Using AERMOD, I estimate that the marginal willingness to pay (MWTP) to reduce exposure to NO_x emissions by 1 $\mu g/m^3$ is \$3,272, an order of magnitude larger than past estimates and more in line with the expected health benefits. This implies that the benefit of the RECLAIM cap-and-trade program to local residents is roughly \$502 million per year, much larger than the estimated abatement costs of \$38 million per year.[2] However, when I use methods now standard in the literature, the estimated price response and the implied social valuation of RECLAIM are small or wrongly signed and not statistically different from zero. AERMOD also makes it possible to address other questions, such as whether people fully respond to invisible pollutants.

Many pollutants are hard to detect without instrumentation, which raises the possibility that buyers with preferences about pollution exposure may suffer from imperfect information or salience effects. I test for this by exploiting the chemical relationship of NO_x and ozone. NO_2, a primary component of NO_x, is highly visible, but under certain atmospheric conditions it transforms into ozone, which is invisible but far more toxic. The conversion rate of NO_x to ozone varies predictably over the course of the year, making it possible to test whether prices depend on long-run expectations based on all currently available information, or whether they are sensitive to foreseeable short-run changes in toxicity and visibility. I find that prices are much more sensitive to visible NO_x than they are to invisible ozone, consistent with a model where buyers suffer from imperfect information, salience effects, or both.

The results suggest that reducing air pollution is a very cost-effective way to improve welfare; however, in Chapter 2, I find that air quality policies may also face a steep trade-off between efficiency and equity. More generally, the failure of standard methods raises the possibility that current estimates of pollution's effects on other outcomes like health and mortality are also significantly too small.

Before presenting the results in detail, I discuss the puzzle of clean air's seemingly low value, what is behind the puzzle, and how to solve it (Section 1.2). I then provide the theoretical framework I use to draw conclusions about MWTP, how pollutant visibility

2. This is the total benefit of decreasing NO_x emissions from actual 1995 levels to the RECLAIM cap in 2005, annualized at 3%. Abatement cost is based on SCAQMD's tabulation of firms' actual equipment and the available abatement technology that would need to be installed to meet certain abatement goals. The cost includes installation and ongoing operation of equipment (SCAQMD 2000).

affects agents' behavior, and how people sort geographically in response to changes in pollution (Section 1.3). Next, I discuss my research design based on the Electricity Crisis and outline my estimation strategy (Section 1.4). Finally, I describe the data I use (Section 1.5), present the results (Section 1.6), and discuss the possible welfare implications (Section 1.7).

1.2 Finding the Value of Clean Air

House prices have long been used to measure the marginal willingness to pay (MWTP) for non-market goods. By varying a single characteristic of a house and observing the associated price change, we can infer the MWTP for that characteristic (see Section 1.3.1). The MWTP for pollution abatement has been measured this way many times, starting with Ridker and Henning (1967).

But past work suggests that house prices do not respond much to pollution, implying a disparity between the MWTP for pollution reductions and the expected health benefits (Smith and Huang 1995). For a family of two adults and one child, 1 microgram of particulate matter ($PM_{2.5}$) per cubic meter of air ($\mu g/m^3$) costs about \$1,600 in increased mortality risk alone, to say nothing of acute illness risk, decreased quality of life, or the costs borne by other family members.[3] However, in their meta-analysis of OLS estimates of MWTP, Smith and Huang (1995) find that the interquartile range of estimated MWTP is \$0 to \$233 per $\mu g/m^3$ TSP and that the mean estimate only covers 6–33% of VSL-based mortality cost.[4]

More recent instrumental variables estimates have not narrowed this disparity. Chay and Greenstone (2005) use the implementation of the National Ambient Air Quality Standards

3. The mortality value for an adult is \$680 and based on the value of a statistical life (VSL) for adults aged 35–44 from Aldy and Viscusi (2008) and adult $PM_{2.5}$ mortality risk from Pope et al. (2002). For a child, the value is \$250 using infant $PM_{2.5}$ risk from Woodruff, Parker, and Schoendorf (2006) and the VSL of a 18–24-year-old, the lowest age estimated by Aldy and Viscusi. All monetary values in the paper are denominated in 2014 dollars unless otherwise noted.

4. There are several measures of the class of pollutants called "particulate matter," which are larger solid and liquid particles rather than gaseous molecules. $PM_{2.5}$ is all such particles with a diameter no larger than 2.5 micrometers (μm), while PM_{10} particles have a diameter between 2.5 and 10 μm. Total suspended particulates, or TSP, is a another measure that corresponds to all particles smaller than 25–40 μm, depending on the apparatus collecting samples. Because of the inconsistent apparatus-dependent definition of TSP, the EPA abandoned it as an official measure in 1987 (52 FR 24634).

(NAAQS), county-level house prices, and average county pollution monitor readings to estimate a MWTP of \$191 for a 1 $\mu g/m^3$ reduction in TSP, well within Smith and Huang's interquartile range. Bayer, Keohane, and Timmins (2009) also use county-level data and use pollution from distant sources as an instrument for local pollution to estimate a MWTP of \$131 per $\mu g/m^3$ reduction of PM_{10}.[5]

This disparity appears to be peculiar to air pollution, as prices readily respond to other location-specific amenities. Cellini, Ferreira, and Rothstein (2010) use house price responses to bond override elections and estimate the average household is willing to spend \$1.50 for a \$1 increase in school capital expenditures. Linden and Rockoff (2008) find that when a registered sex offender moves into a neighborhood, the value of nearby houses drops by about \$7,000, more than the FBI's estimates of victimization costs would suggest. Davis (2004) looks at how prices respond to the appearance of a cancer cluster in Churchill County, Nevada, where the rate of pediatric leukemia suddenly skyrocketed for unknown reasons. The price response there implies the welfare cost of pediatric leukemia is about \$7 million, in line with estimates of the value of a statistical life from Aldy and Viscusi (2008).

The disparity is also not caused by a general ignorance of pollution's health costs or an unwillingness to avoid pollution. For example, it could be that people simply do not know that pollution is dangerous, or that, like junk food, the cost of a marginal dose is not salient enough to elicit a behavioral response. However, Neidell (2009) and Moretti and Neidell (2011) find the opposite. They find that attendance at outdoor attractions like the zoo and sporting events drops precipitously in response to smog alerts, suggesting that people not only know the health risks but are willing to undertake costly avoidance behavior.

This body of conflicting evidence suggests that something specific to air pollution is attenuating house price responses or estimates of those responses.

---

5. The estimate from Chay and Greenstone (2005) is based on their preferred specification in Table 5A, column 4. The estimate from Bayer, Keohane, and Timmins (2009) is taken from Table 6, column 2. This estimate assumes costless migration, which is standard in the literature. They also fit a structural model that allows for costly migration, which yields a MWTP estimate of \$352.

1.2.1 Econometric Problems Behind the Puzzle

A likely candidate for attenuation bias is misspecification in who is exposed to pollution (or pollution clean-up) and who is not. This is because, unlike wages or education, there are no data on individual-level pollution exposure, so researchers must approximate exposure in some way. In the economics literature, two approaches are predominantly used.[6] The first and most straightforward approach is to use a geographic difference-in-differences design where people close to a pollution source are assumed to be exposed to the source while those slightly farther away are assumed not to be exposed. The second approach is to use data from the EPA's network of pollution monitors as a proxy for person-, neighborhood-, or county-level exposure, usually by interpolating between monitors.

Unfortunately, both of these methods suffer from the same problem: They are unable to capture sharp changes in pollution exposure across short distances, which biases estimates based on these methods. It is also important to note that these problems are inherent to pollution exposure generally and thus extend to estimates of pollution's effect on any outcome.

Bias in Geographic Diff-in-diff Estimates

In a geographic difference-in-differences design, people around a pollution source are assigned to treatment and control groups based on their proximity to the source. The econometrician chooses radius r_0 around the source to define the treatment group and radius $r_1 > r_0$ to define the control. Having defined treatment and control groups, the problem is now a standard diff-in-diff around some shock to the source's pollution emissions. This allows the reduced-form effect of the pollution source to be estimated when data on exposure is unavailable.[7]

When used to study air pollution, however, the geographic diff-in-diff is biased because

6. Currie et al. (2014) summarize the methods used in the literature on pollution's effect on children's health.

7. For examples of research focused on reduced-form geographic diff-in-diffs, see Currie and Walker (2011) and Currie et al. (2015).

the wind does not respect the radii chosen by the econometrician and contaminates both the treatment group and the control group. Suppose the true effect of a polluting firm on outcome y_{it} is

$$y_{it} = \alpha N_{it} + \beta X_{it} + \varepsilon_{it} \tag{1.1}$$

where X_{it} is pollution exposure to i at time t, N_{it} is exposure to other disamenities created by the firm (e.g., eyesore of a refinery), and $t \in \{0, 1\}$ indexes the pre- and post-shock time periods, respectively. Exposure can be written $X_{it} = m_{ft} \cdot h(r_{fi}, \theta_{fi}; \boldsymbol{S}_f)$ where m_{ft} is firm f's emissions and h is the probability density function that a molecule of emissions ends up at distance r and heading θ relative to the firm. The vector $\boldsymbol{S}_f$ contains variables about the physical characteristics of the firm's polluting equipment (e.g., height of the smoke stack) and local meteorological conditions like wind speed and direction. Assume r_0 is chosen so that $r_{fi} > r_0$ implies $N_{it} = 0$.

The geographic diff-in-diff estimates the reduced form as

$$y_{it} = \gamma_1 + \gamma_2 \cdot \text{post}_t + \gamma_3 \cdot C_i + \gamma_{\text{GD}} \cdot \left(C_i \times \text{post}_t\right) + \mu_{it} \tag{1.2}$$

where $C_i = \mathbf{1}\{r_{if} \leq r_0\}$ is a dummy variable for individuals living in the treatment area. We can write the expected value of y_{it}, conditional on i's treatment assignment, in terms of the average effects on the treatment group:

$$\mathbb{E}_i[y_{it} \mid C] = \alpha \bar{N}_t^C \cdot C + \beta \bar{X}_t^C \cdot \left[C + \varphi(1 - C)\right] \tag{1.3}$$

where $\bar{X}_t^C = \mathbb{E}_i[X_{it} \mid C = 1]$ and $\varphi = \mathbb{E}_i[X_{it} \mid C = 0] / \bar{X}_t^C$. Figure 1.1 depicts the geographic diff-in-diff's radii with the true downwind treatment marked by the shaded region and the average effect for each area based on Equation (1.1). Note that as wind speed increases, the shaded treatment area narrows and extends farther from the source, increasing relative exposure downwind and thus increasing φ.

By construction, the geographic diff-in-diff recovers the following estimate of pollution's

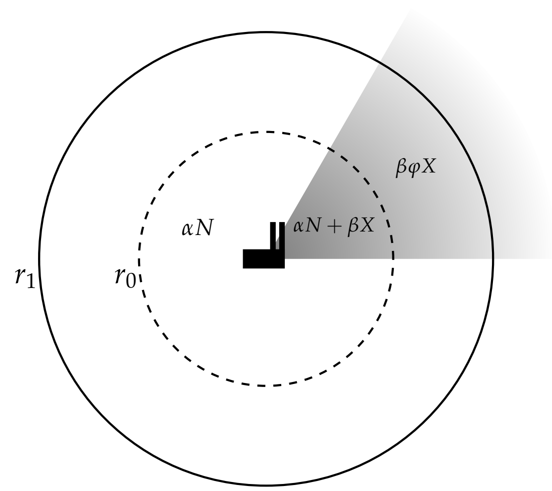

Figure 1.1: *Geographic Diff-in-diff with Wind*

Notes: Dashed circle denotes boundary of geographic diff-in-diff's treatment group, solid circle denotes boundary of control group and sample. Shaded area is the true treatment area downwind. Values r_0 and r_1 are treatment and control radii. Other values are reduced form effects of the firm, see Equation (1.1) and Section 1.2.1.

effect on y_{it}:

$$\begin{aligned}\hat{\gamma}_{\text{GD}} = \quad & \mathbb{E}[y_{it} \mid C = 1, \text{post} = 1] - \mathbb{E}[y_{it} \mid C = 1, \text{post} = 0] \\ & - \left(\mathbb{E}[y_{it} \mid C = 0, \text{post} = 1] - \mathbb{E}[y_{it} \mid C = 0, \text{post} = 0]\right)\end{aligned}$$

Using Equation (1.3), this reduces to

$$\hat{\gamma}_{\text{GD}} = \underbrace{\alpha\left(\bar{N}_1^C - \bar{N}_0^C\right)}_{\text{Non-pollution Effect}} + \underbrace{(1-\varphi)}_{\text{Wind bias}} \cdot \underbrace{\beta\left(\bar{X}_1^C - \bar{X}_0^C\right)}_{\text{Pollution Effect}} \tag{1.4}$$

The first term captures the firm's non-pollution effects. As β is the coefficient of interest, the ideal research design would hold N_{it} constant over time, making this term 0.[8] The second term is the change in average exposure to the treatment group, multiplied by the contamination factor $(1 - \varphi)$.

Thus, even when non-pollution effects are held constant over time, the estimate of the pollution effects is biased because the control group is actually treated as well. And because φ increases with wind speed, the contamination factor $(1 - \varphi)$ and $\hat{\gamma}_{\text{GD}}$ both become more negative as wind speed increases. Furthermore, because the distribution function h need not be monotonic in r, φ need not be less than 1, meaning $\hat{\gamma}_{\text{GD}}$ could have the wrong sign.[9] This contamination problem is common in program evaluation (e.g., Miguel and Kremer 2004) and can be fixed by re-scaling by average treatment intensity. But this requires a good measure of treatment intensity, and, as Section 1.2.1 argues, this is a role monitor data are not well suited to play.

Empirically, the dependence of the bias on wind speed is important for two other reasons. First, greater Los Angeles is one of the least windy areas in the United States, so if the wind significantly biases estimates in this sample, it almost certainly biases estimates in other regions with greater wind speeds. Second, when pollution is less influenced by the wind,

8. This is naturally not the case when the shock to the firm is the construction of the firm itself (as in Banzhaf and Walsh 2008, Davis 2011, and Currie et al. 2015). In such cases, $\bar{N}_1^C > \bar{N}_0^C = 0$. Note also that as the wind gets stronger and $\varphi \to 1$, $\hat{\gamma}_{\text{GD}} \to \alpha\bar{N}_1^C$ and the geographic diff-in-diff recovers the *non*-pollution effects of the firm, including sorting effects for outcomes not directly impacted by non-toxic disamenities.

9. An example of the non-monotonicity of exposure with distance is given by Figure 1.2b.

standard non–wind-based estimates should be less biased. For example, when pollution is emitted at ground level, more of it stays close to the source, keeping φ low. This suggests that the bias in geographic diff-in-diffs around pollution from vehicles on the ground (e.g., Currie and Walker 2011) may not be as severe. However, even car exhaust gets carried by the wind (Hu et al. 2009), and a low φ does not mitigate the separate bias introduced by monitor data.

Bias from Pollution Monitor Interpolation

The most common way monitor data are adapted for use in economics is interpolation, which uses data from pollution monitors to approximate pollution exposure at other locations of interest.[10] To study the effects of pollution exposure on some outcome, we need data on the outcome and pollution exposure, $\{(y_i, x_i)\}_{i=1}^{N}$, but x_i is never observed. However, we do observe $\{x_m\}_{m=1}^{M}$, pollution exposure at monitor locations. If values of x are spatially correlated, so $\text{Cov}(x_i, x_j)$ is high if i and j are physically close to one another, then $\{x_m\}$ can be used to construct an approximation $\tilde{x}_i$ for any needed x_i.

The viability of any interpolation method depends critically on the spatial covariance of x. In the most extreme case where $\text{Cov}(x_j, x_k) = 0$ for all $j \neq k$, the interpolated values will obviously be no better than random noise because the monitor data $\{x_m\}$ do not provide any information about x_i. Similarly, if $\text{Cov}(x_j, x_k)$ falls quickly as the distance between j and k grows, then more monitors will be needed at a higher spatial frequency to cover the sample area. For example, if $\text{Cov}(x_j, x_k) \approx 0$ if j and k are more than 1 km apart, but all monitors are 5 km apart, then the interpolated $\tilde{x}_i$ will be no better than noise for large portions of the sample area. The converse also holds and helps explain why rainfall data, which is highly correlated across tens of kilometers, has been successfully interpolated in

10. It is also possible to use monitor data by restricting the data sample to people living close to a single monitor. The shortcomings of this method are entirely practical, since reducing the sample radius reduces the measurement error but also reduces the sample itself. This method works well in case studies, like Graff Zivin and Neidell's (2012) analysis of how worker productivity at a single firm covaries with readings from a nearby $PM_{2.5}$ monitor. However, the trade off between sample size and measurement error limits its large-scale use.

many different contexts.[11]

Unfortunately, air pollution exhibits a much lower degree of correlation across space because of the discrete nature of pollution sources. Unlike rainfall and other continental-scale geologic processes, air pollution is predominantly created by many distinct sources like firms and cars. This makes the local geographic distribution of pollution exposure very idiosyncratic, with sharp changes over very short distances; e.g., pollution levels downwind of a highway are dramatically different from pollution levels upwind. This in turn means that the relationship between any given x_i and a monitor reading x_m depend on many more factors than distance and relative direction. Most importantly, $\mathrm{Cov}(x_i, x_m)$ depends on whether a major pollution source exists between i and m. If m is downwind of the source, x_m varies with the source's emissions but x_i does not, and vice versa.

Evidence confirming this problem can be found in existing literature, even though the problem itself has not been directly raised or addressed. Studies using interpolated values often present a leave-one-out cross-validation as evidence of the interpolation's quality.[12] The value of each monitor reading x_m is interpolated using all remaining monitors and the correlation between x_m and $\tilde{x}_m$ is calculated, with a high correlation coefficient assumed to be evidence of a good interpolation. However, the correlation of x_m and $\tilde{x}_m$ presented in these studies is generally unconditional, which conflates spatial correlation with secular temporal correlation which may equally effect all monitors (e.g., seasonal trends in ozone). Karlsson, Schmitt, and Ziebarth (2015) use German pollution monitors and inverse distance weighting (IDW) to calculate this cross-validation correlation for several pollutants in Germany. The unconditional correlations range from 0.5 to 0.93; however, conditional on year and season effects, the correlations drop precipitously, ranging from 0.15 to 0.47.[13] Likewise, Knittel, Miller, and Sanders (2014) and Lleras-Muney (2010) present

11. See Pouliot (2015) for a summary of rainfall interpolations.

12. Inverse distance weighting, as well as this cross-validation technique, has been the standard method for with sub-county pollution analyses in the economics literature since Neidell (2004) and Currie and Neidell (2005).

13. See Table F1 of Karlsson, Schmitt, and Ziebarth (2015).

evidence of non-classical measurement error in IDW and Kriging interpolations, respectively. In this context, the non-classical measurement error will exacerbate the usual attenuation in OLS estimates from classical and potentially cause wrongly signed estimates.[14]

And unlike classical measurement error, quasi-experimental research designs and other IV methods will not necessarily redeem a bad interpolation. Given an instrument z and interpolation error $\eta = \tilde{x} - x$, estimates will only be consistent if z and η must be uncorrelated. This simply follows from the canonical probability limit of the IV estimator:

$$\operatorname{plim} \hat{\beta}_{\text{IV}} = \beta \cdot \frac{\text{Cov}(x,z)}{\text{Cov}(x,z) + \text{Cov}(\eta,z)} \tag{1.5}$$

Note that in this case that $\hat{\beta}_{\text{IV}}$ could be bigger or smaller than β depending on the joint distribution of (x, z, η) which will vary across research designs. Nevertheless, $\hat{\beta}_{\text{IV}}$ can only be consistent when $\text{Cov}(\eta, z) = 0$, which is unlikely to be the case in the most commonly used research designs.

In the case of the geographic diff-in-diff, this condition is very unlikely to hold because firms outnumber monitors by several orders of magnitude. According to the EPA's AirData summary files, the average county had 1.01 monitors in 2005, with almost two-thirds of counties having zero monitors. Despite being in one of the most intensively studied areas in the United States, each monitor in greater Los Angeles is outnumbered by hundreds of firms. This disparity is readily apparent in Figure A.2, which maps the locations of every polluting firm and pollution monitor in the area. With so few monitors, the distribution of $\tilde{x}$ will be smooth across the sample area of most firms; that is, $\tilde{x}$ will not spike downwind of the firm. But actual exposure x does spike, particularly close to the firm, so η will also spike near the firm and will be correlated with proximity to the firm. And since the instrument z is defined by proximity to the firm, $\text{Cov}(\eta, z) \neq 0$.

In the case of county-level studies using the Clean Air Act (CAA) as a natural experiment, estimates are also likely to be inconsistent because z is mechanically related to $\tilde{x}$. The CAA

14. This is because the distribution of $\tilde{x}$ is smoother than that of x, so $\text{Var}(\tilde{x}) < \text{Var}(x)$. Noting that $\tilde{x} = x + \eta$, where η is the interpolation/measurement error, it immediately follows that $\text{Cov}(x, \eta) < 0$.

is often used as a natural experiment because it instituted more stringent regulations on counties whose average monitor reading exceeded a certain threshold. In these county-level studies, the measure of pollution exposure to the county, $\tilde{x}$, is generally the very same monitor average that affects a county's treatment status. The econometric problem this causes is easier to see by noting that Equation (1.5) can also be written as

$$\operatorname{plim} \hat{\beta}_{\mathrm{IV}} = \beta \cdot \frac{\operatorname{Cov}(x, z)}{\operatorname{Cov}(\tilde{x}, z)} \tag{1.5$'$}$$

Thus, if the treatment impacts monitor readings $\tilde{x}$ more than actual exposure x, $\operatorname{plim} \hat{\beta}_{\mathrm{IV}} < \beta$ and the estimate will understate the true effect. This would be the case if, as Bento, Freedman, and Lang (2015) find, regulators put more effort into reducing pollution levels at problematic monitors within the county.[15]

1.2.2 Solving the Puzzle with Atmospheric Dispersion Modeling

The econometric problems described above are rooted in the idiosyncratic nature of pollution exposure across space. Any measure of pollution must capture sudden changes in exposure over short distances in order to be useful in statistical analyses. Atmospheric dispersion models use detailed data on meteorology and firms to accomplish exactly this goal.

A dispersion model uses data on a firm's polluting equipment and the meteorology around the firm (the vector $\boldsymbol{S}_f$ from Section 1.2.1) and predicts the spatial distribution of the firm's pollution (the function h from Section 1.2.1). In this paper, I use AERMOD, the EPA's legally preferred model for short-range applications. This preference is based on the model's high accuracy as established by peer-reviewed field tests (e.g., Perry et al. 2005).[16] To account

15. There is a more general problem with using the average of a county's monitors: The relationship between $\tilde{x}$ and the true distribution of individual-level exposure is unclear and changes over time because monitors are a sample across *space*, not population. Even if it could be credibly established that $\tilde{x}$ is an unbiased approximation of the mean (or any order statistic) of the true exposure distribution at some point in time, this relationship would quickly be broken as people and firms change their behavior and locations over time.

16. Regulatory preference is stated in 40 CFR pt. 51, app. W (2004). See Cimorelli et al. (2005) for a rigorous development of the model itself. Field tests are generally conducted by placing several dozen monitors around a polluter and adding to its emissions a non-toxic, non-reactive tracer chemical which is not usually present in the area.

for meteorological conditions, AERMOD requires hourly data on temperature, wind speed, and wind direction at multiple elevations; the standard deviation of vertical wind speed; the convectively and mechanically driven mixing heights; and other parameters.[17] AERMOD also requires five parameters for the pollution source itself: the smoke stack's height and diameter, the temperature and velocity of the gas exiting the stack, and how much pollution is emitted by the stack.

Using these data, AERMOD yields $\text{aermod}_{ift} = \text{NOx}_{ft} \cdot \tilde{h}(r_{fi}, \theta_{fi}; \boldsymbol{S}_f)$, the pollution exposure to location i at time t due to NO_x emissions from firm f, measured in micrograms per cubic meter of air (µg/m^3). Summing over all firms yields total industrial exposure: $\text{aermod}_{it} = \sum_f \text{aermod}_{ift}$. It is important to note that these AERMOD-based measures do not represent NO_x exposure alone. AERMOD uses data on how much NO_x a firm emits, but NO_x will react in the atmosphere to become ozone, thus aermod_{ift} is a composite measure of NO_x and ozone exposure. Section 1.4.2 below describes how this fact can be used to test how buyers respond differentially to NO_x, which is visible, and ozone, which is not.

Mapping aermod_{ift} (Figure 1.2) and aermod_{it} (Figure 1.3) makes clear the problems caused by geographic diff-in-diffs and monitor interpolation, respectively. Figure 1.2 maps aermod_{ift} for a single firm, the Scatterwood Generating Station. The concentration of NO_x-based pollution is plotted for all 100-meter grid squares. For Figure 1.2a, this is limited to area less than 20 kilometers of the firm. Figure 1.3a maps aermod_{it}, total exposure to industrial emissions, across the entire sample area, with monitor locations marked by white dots.[18] Figure 1.3b shows how exposure would be calculated by interpolating aermod_{it} from actual monitor locations.

Figure 1.2 shows that the direction and speed of the wind is crucial in knowing who is affected by the firm. It also shows how extensive the contamination of a geographic diff-in-diff can be. In particular, Figure 1.2b offers a closer look at the exposure around

17. A full list of the variables used is found in the AERMOD user manual or Cimorelli et al. (2005).

18. Section 1.4.3 discusses how this sample region is defined. Details about how AERMOD is implemented in this paper are given in Section 1.5.4.

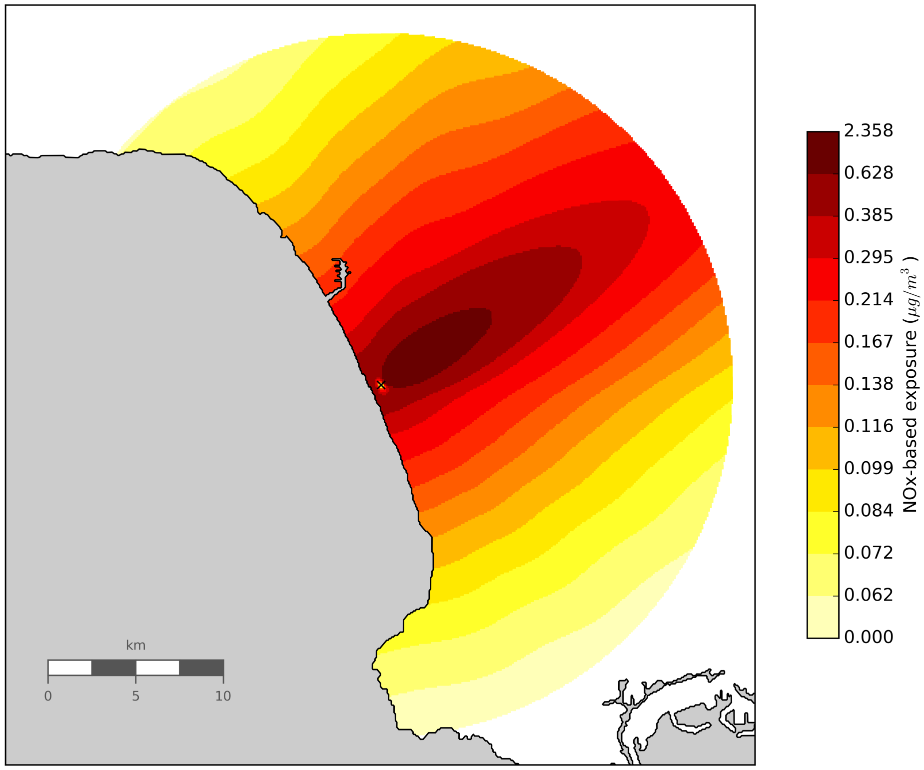

(a) *All land within 20 km*

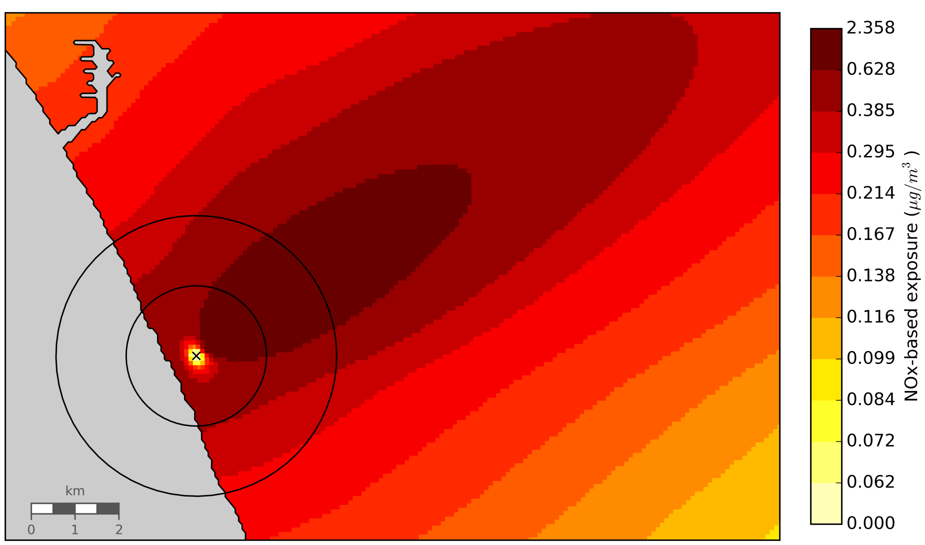

(b) *Zoomed*

Figure 1.2: *Exposure due to Scatterwood Generating Station, 1999*

Notes: The color of each square is determined by the average aermod exposure due to the Scatterwood Generating Station in 1999. Each plotted square is 100 meters wide. Breaks in the color scale are set at order statistics of the plotted sample: minimum, 1–9th deciles, the 95th and 99th percentiles, and the maximum.

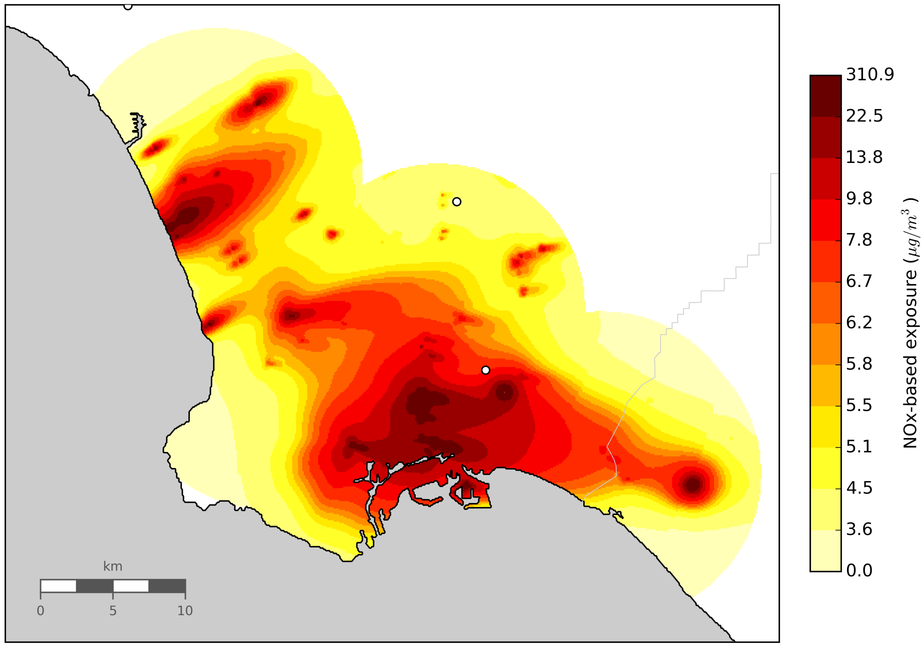

(a) *AERMOD-based Exposure*

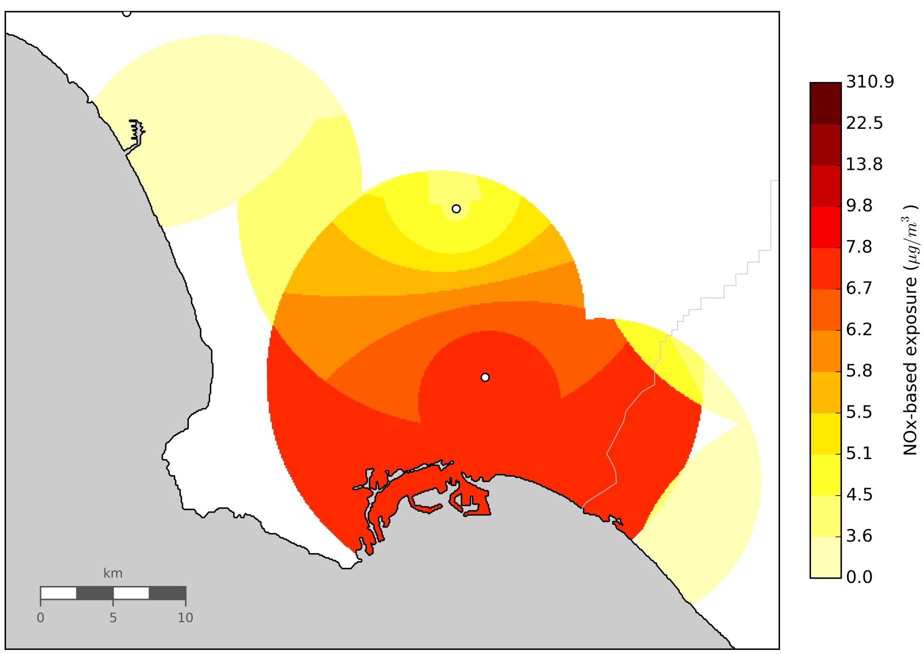

(b) *Interpolated (IDW) AERMOD-based Exposure from Monitor Locations*

Figure 1.3: *Total Exposure in Sample Region, 1999*

Notes: Sub-figure (a) plots the average exposure due to all firms. Sub-figure (b) plots the average exposure as interpolated from monitor locations marked with black circles. Breaks in the color scale are set at order statistics of the plotted sample in sub-figure (a): minimum, 1–9th deciles, the 95th and 99th percentiles, and the maximum. Each plotted square is 100 meters wide.

the firm, with circles drawn at one and two miles from the firm for easy comparison to a geographic diff-in-diff's treatment and control groups. Much of the control group sees extreme levels of exposure while the area of lowest exposure in the geo diff-in-diff sample is actually in the treatment group.[19]

Figure 1.3 shows that there is far too much spatial variation in exposure to be captured by so few monitors. Figure 1.3a shows how quickly exposure can change over short distances and how unpredictable the exposure distribution can be. The number of local extrema and inflection points far exceeds the number of nearby monitors. Figure 1.3b makes this problem easier to see by showing the interpolated values of aermod_{it} based on the actual monitor locations. The interpolation follows the literature and is calculated using inverse distance weighting (IDW) with monitors restricted to those with full NO_x coverage over the sample period (1997–2005). Monitors are also given a weight of zero if they are more than 15 km from the point of interest.

Little of the variation seen in Figure 1.3a remains after interpolation. Most locations' predicted exposure are perfectly correlated with the nearest monitor, and the area that does have some variation at best vaguely resembles the true distribution. Note that if the 15-km interpolation radius were expanded, this would add no true variation to the data because the sample of monitors would be the same. This would be especially troublesome for the southwestern corner of the region, the Palos Verdes Peninsula, which has very low exposure because it is upwind of all major polluters. Despite being one of the cleanest areas in the sample, it would be assigned a very high-level exposure and be indistinguishable from the truly polluted area near the monitor.

The complex patterns seen in the wind-based exposure distribution is obviously difficult to approximate using concentric circles or other simple methods. The possibility that factors like the wind might affect estimates has been raised occasionally in the literature, but the results have not suggested it is an important issue. Of the economics papers on industrial

19. This non-monotonicity is caused by the height of the firm's smoke stacks (about 300 feet) and the buoyancy of the hot gases they emit. The bulk of the smoke plume travels laterally in the air before touching down.

pollution that have tried to account for the wind, only Hanna and Oliva (2015) find that the wind significantly alters their estimates, and then only in certain specifications.[20] The mixed nature of these past results is likely due to the complexity of the atmospheric dispersion problem, which has been the dedicated focus of many atmospheric scientists for decades (see Cimorelli et al. 2005, Section 1 for a summary). Fortunately, the econometric problems described above can be avoided by taking advantage of their work.

1.3 Theory and Predictions

This section presents a simple model of locational choice and describes how it can be used to answer the economic questions of interest: what are people willing to pay for clean air; does the market fully capitalize the costs of invisible pollution; and what is the incidence of an air quality improvement.

1.3.1 House Prices, Hedonics, and MWTP

When choosing a place to live, households weigh a location's amenities and house prices against their own income and preferences. They solve

$$\max_{c,\boldsymbol{g}} u(c, \boldsymbol{g}; \boldsymbol{\alpha}) \quad \text{s.t.} \quad y = c + P(\boldsymbol{g}) \tag{1.6}$$

where c, the numeraire, is aggregate non-amenity consumption; $\boldsymbol{g}$ is a vector of public and private amenities provided by the chosen neighborhood and house; $P(\boldsymbol{g})$ is the price of a house with amenities $\boldsymbol{g}$; and $(y, \boldsymbol{\alpha})$ are income and a vector of preference parameters, respectively, and together define the household. This differs from a standard consumer problem because many elements of $\boldsymbol{g}$, like air quality or proximity to the ocean, are location

20. Hanna and Oliva (2015) look at how labor supply in Mexico City responded to a drop in pollution after the closure of a large refinery. They include the local elevation and a linear measure of degrees downwind in some specifications. Davis (2011) estimates the effect of plant openings on nearby house values and includes dummy variables for "upwind" and "downwind" in a robustness check. Schlenker and Walker (Forthcoming) measure the change in daily hospital visits due to changes in airport traffic and incorporate wind speed and direction into one of their models. Luechinger (2014) compares county-level infant health before and after the mandated desulfurization of power plants in Germany. He calls a county "downwind" of the power plant if it falls in the same 30-degree arc as the prevailing wind direction and includes downwind dummies in all his specifications.

specific, so households must physically relocate in order to change their consumption of these amenities. This adds a spatial element to the standard market clearing equilibrium conditions—every household must weakly prefer their current location to all others.

Rosen (1974) noted that utility-maximizing agents will choose a bundle of amenities and prices $(P(\boldsymbol{g}^*), \boldsymbol{g}^*)$ so that their marginal willingness to pay for each $g_k \in \boldsymbol{g}$ is equal to the marginal price.[21] To see why this is the case, note that for some fixed utility level $\bar{u}$, the solution to Equation (1.6) can be re-written

$$u\left(y - \theta(\boldsymbol{g}^*; y, \boldsymbol{\alpha}, \bar{u}), \boldsymbol{g}^*; \boldsymbol{\alpha}\right) = \bar{u} \tag{1.7}$$

where θ is the agent's willingness to pay for $\boldsymbol{g}$, conditional on $(y, \boldsymbol{\alpha}, \bar{u})$. For a single amenity g_k, $\partial\theta/\partial g_k = \theta_{g_k}$ is the marginal willingness to pay for g_k, and P_{g_k} is the marginal price for g_k. If $\theta_{g_k} > P_{g_k}$, then the agent can buy more g_k for less than she would otherwise be willing to pay, and vice versa if $\theta_{g_k} < P_{g_k}$; thus in equilibrium, $\theta_{g_k} = P_{g_k}$ for all g_k at $\boldsymbol{g}^*$.

Estimating the average MWTP, which is difficult to do directly, can therefore be accomplished by estimating P_{g_k} instead, though this requires some assumptions. In order to identify P_{g_k} using intertemporal variation in house prices, the shape of P, which is endogenously determined in equilibrium, must be constant over the sample period (Kuminoff and Pope 2014). While this assumption is less palatable for longer sample periods and low-frequency data, it is likely to hold when using a short sample period and quarterly data. Another potential problem is that $(P(\boldsymbol{g}^*), \boldsymbol{g}^*)$ is endogenously chosen by the agent, creating a potentially omitted variables problem (Bartik 1987; Epple 1987; Chay and Greenstone 2005). Any attempt to identify P_{g_k} must address this and satisfy the identification assumptions specific to the chosen research design, which I discuss for this paper in Section 1.4.

21. There are a number of theoretical frameworks that can be used to estimate MWTP. See Palmquist (2005) and Kuminoff, Smith, and Timmins (2013) for summaries of valuation using hedonic pricing and equilibrium sorting models.

1.3.2 Pollutant Visibility and Prices

Even if people care about pollution, they cannot bid more for houses with cleaner air if they cannot discern clean air from dirty air.

In an efficient market, a house's price should reflect the net present value of expected future utility flows afforded by the house's amenities. If amenities change—or people believe they will change—it should be reflected immediately in the market price of the house. Therefore, any transitory or already foreseen changes in pollution levels, like predictable seasonal variation, should not affect a house's price.

Conversely, if buyers suffer from imperfect information or salience effects, then prices may depend on transitory changes in pollutant concentrations or salience. Determining current pollution levels is difficult without equipment because many pollutants, like ozone, are colorless and have no bad smell. Extrapolating pollution's daily, weekly, and yearly patterns after a single viewing is even more difficult. But even with perfect information, people may not respond to pollution if it or its costs are not salient. There is a growing body of evidence that salience and framing can significantly affect even weighty decisions like choosing a house and a neighborhood.[22]

We can distinguish between these cases empirically by testing whether house prices respond to foreseeable changes in the composition of air pollution. NO_x is emitted directly by polluters and becomes ozone at different but predictable rates throughout the year. Thus, with perfect information and rational agents, house prices should not respond to these seasonal changes. If the price response does vary seasonally, the physical characteristics of NO_x and ozone will allow us to identify whether toxicity or visibility affects buyers more (see Section 1.4.2).

22. Pope, Pope, and Sydnor (2014) show that house prices gravitate toward round numbers like $150,000, suggesting that psychological biases play a large role in major purchases. Busse et al. (2015) find people are more likely to buy a convertible car on a hot or cloud-free day, even if they have already owned a convertible and should know how much utility they get from driving a convertible in the snow. An earlier version of this paper, Busse et al. (2012), also provided evidence that houses with air conditioners and swimming pools fetch higher prices during the summer.

1.4 Research Design

In this section, I describe how I use the California Electricity Crisis as a natural experiment (Section 1.4.1) and how the Crisis shocked both NO_x and ozone, which can be used to identify the affect of pollutant visibility on prices (Section 1.4.2). Section 1.4.3 details the econometric models I estimate.

1.4.1 Electricity Crisis as Natural Experiment

Estimates of pollution's effect on house prices may suffer from omitted variables bias because households endogenously choose their bundle of amenities and many characteristics about the location and the residents themselves are unobservable. To identify the causal effect of pollution exposure on house prices, I use the natural experiment created by the California Electricity Crisis of 2000, which unexpectedly and permanently lowered NO_x emissions through its effect on the RECLAIM cap-and-trade program.

In 1994, the South Coast Air Quality Management District (SCAQMD), which regulates air pollution in Los Angeles, Orange, San Bernardino, and Riverside Counties, instituted a cap-and-trade program for NO_x emissions called RECLAIM (see Fowlie, Holland, and Mansur 2012). At that time, firms were given an initial allocation of RECLAIM Trading Credits (RTCs) which were tied to a specific year. At the end of each year, firms must surrender one current-year RTC for every pound of NO_x emitted. Excess RTCs can be sold to other firms but not banked for future years. To ease firms' transition into the program, the total number of RTCs was set to be higher than total emissions initially and decrease over time, eventually creating a binding cap.

However, the California Electricity Crisis caused the aggregate cap to bind suddenly, which in turn caused firms to suddenly cut their emissions. Through 1999, most firms had more than enough RTCs to cover their emissions, so there was little need to trade or install abatement equipment. Because of the lack of demand, RTC prices were low, and firms expected that they would be able to buy RTCs cheaply when their own private cap became binding. In 2000, demand for electricity unexpectedly outstripped potential supply,

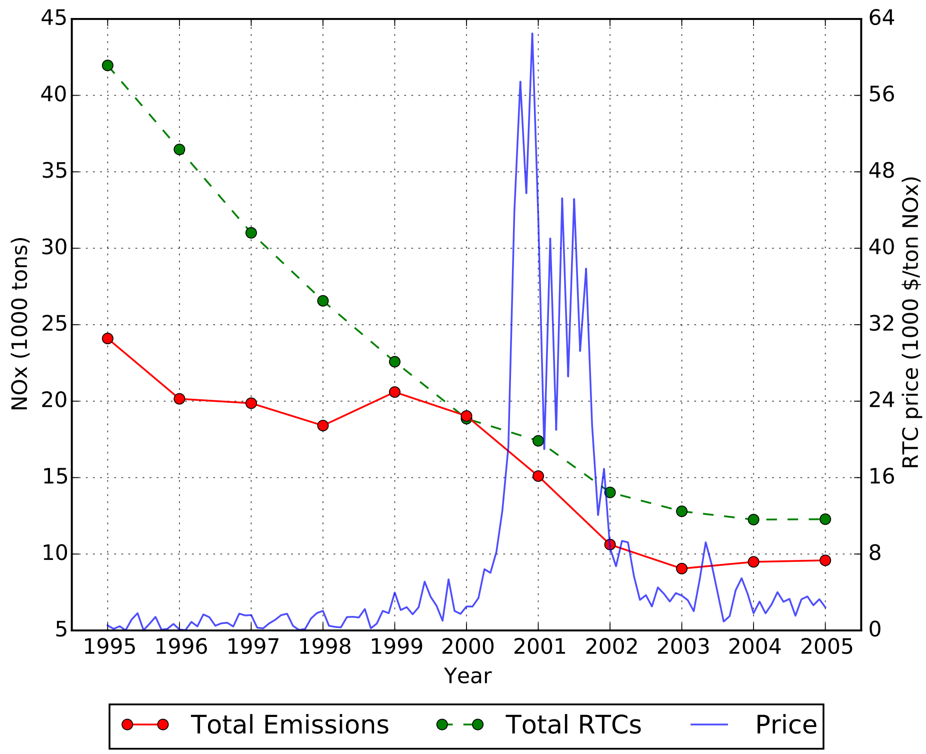

Figure 1.4: *RECLAIM Market*

Notes: "Total RTCs" is the number of RTCs expiring in the calendar year. "Price" is the average of all arms-length transactions in a month across all RTC vintages.

beginning the California Electricity Crisis.[23] Electricity generators ramped up production to prevent rolling black outs. However, generators subject to RECLAIM needed RTCs to cover their increased emissions, which caused the aggregate RTC cap to suddenly bind. RTC prices skyrocketed and non-electric firms cut their emissions in response.

This dramatic change is shown in Figure 1.4, which plots total NO_x emissions, available RTCs for each year, and monthly RTC prices. With the onset of the Crisis, RTC prices jumped from an average of $2,800 in 1999 to a peak of $62,000 at the end of 2000. The resulting drop in emissions is shown in Figure 1.5, which plots the average of firm emissions by quarter and year, giving firms equal weight by re-scaling a firm's emissions by its own sample maximum. Electric generation firms ramped up emissions somewhat in late 1999

23. The exact causes of the shortage and the Crisis in general are a source of much debate. See Borenstein (2002) and Weare (2003), especially Section 3.

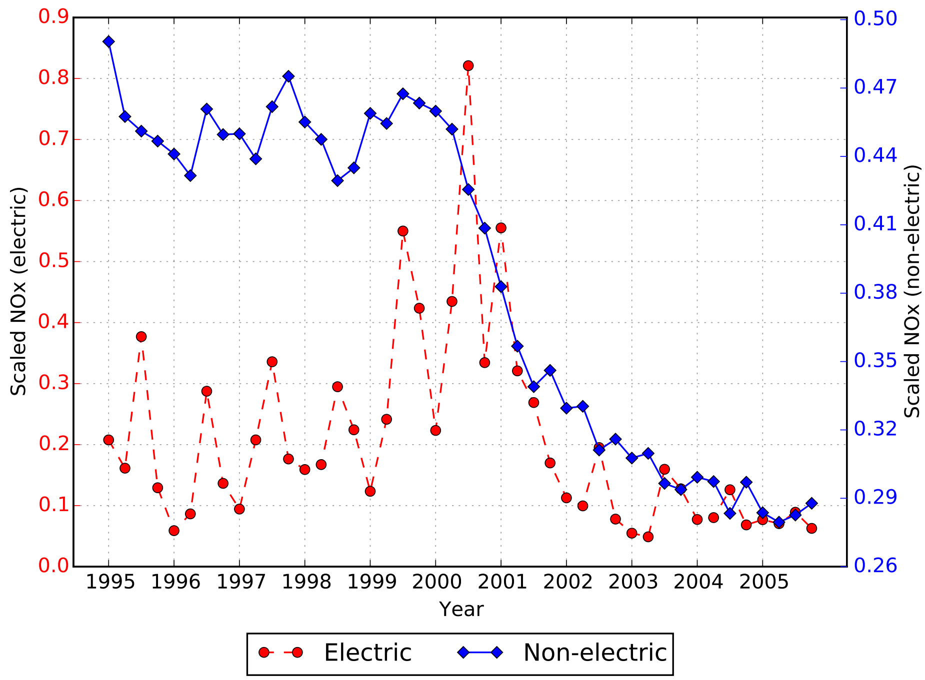

(a) *Quarterly*

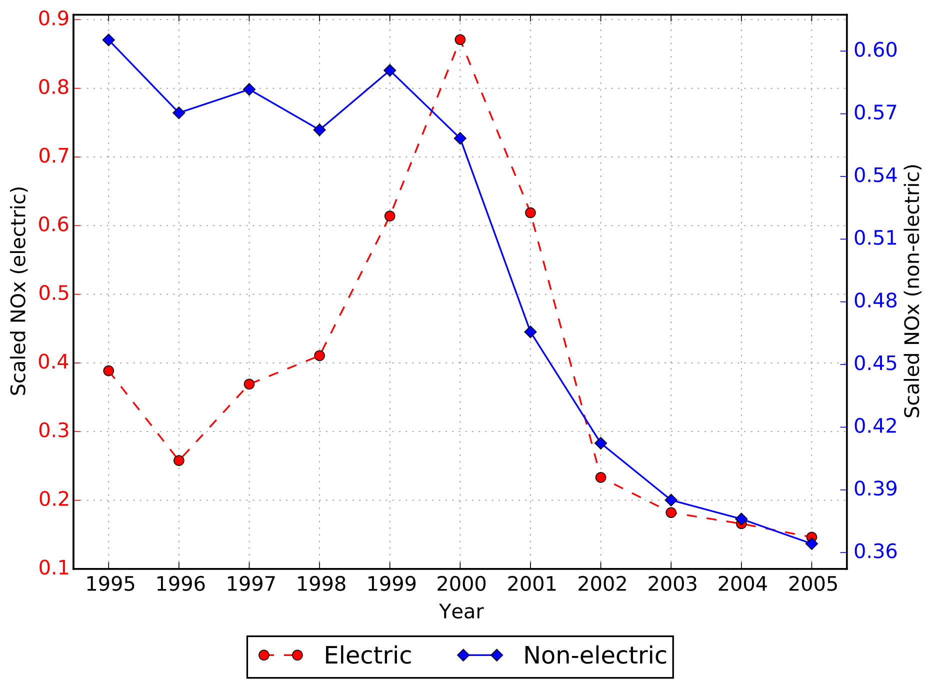

(b) *Annual*

Figure 1.5: *Scaled Emissions by Firm Type*

Notes: Firm emissions are scaled by firm's own maximum emissions. Sample is restricted to firms that operated in at least 8 quarters.

and then in earnest in 2000. Non-electric firms responded by cutting emissions dramatically from the third quarter of 2000 through 2001, with a more modest decline afterwards.

In effect, the Crisis hastened firms' long-run adaptation to a binding cap, causing a sudden and permanent drop in emissions. In general, a firm can reduce its emissions by either lowering production or altering the production process itself, usually by installing equipment which removes NO_x from its combustion exhaust before it reaches the outside air. And while the Crisis was temporary, RECLAIM's binding cap was not, meaning firms had a strong incentive to make long-term adjustments. This is why the temporary Crisis caused the permanent drop in pollution seen in Figures 1.4 and 1.5.

This sudden, permanent drop in emissions can be used to construct a set of instruments for local residents' exposure to firms' pollution. When faced with high RTC prices, firms with more emissions had a larger incentive to cut emissions, so the Crisis should have had a larger effect on houses downwind of these firms. We can use a house's pre-Crisis exposure to gauge how the Crisis changed its exposure relative to other houses. Using aermod_{it}, the AERMOD-predicted exposure to house i in time t, I define pre-Crisis exposure aermod_pre_i as the average exposure across all 8 quarters in 1995 and 1996. The interaction of aermod_pre_i and δ_y, a dummy variable for year y, captures the differential effect of the Crisis on house i in year y. The full set of these interactions $\text{aermod\_pre}_i \times \delta_y$, which I will refer to as the "annual" set of instruments, captures the differential effect of the Crisis on exposure across space and over time.

Similarly, a single interaction, $\text{post}_t = \mathbf{1}\{y \geq 2001\}$, can be used to form a single instrument that I will refer to as the "post" instrument. This instrument, $\text{aermod\_pre}_i \times \text{post}_t$, is the equivalent of a difference-in-difference estimate with variable treatment intensity. While it is coarser than the set of annual instruments, it allows us to summarize the reduced form and first stage effects of the Crisis conveniently with a single number.

The critical identification assumption behind these instruments is that there are no contemporaneous changes in house prices or non-industrial pollution exposure that are correlated with the instruments, conditional on the other covariates. For example, the housing bubble might have induced more appreciation in poorer neighborhoods which also might have been more affected by the Crisis. Fortunately, we can explicitly control for time trends in such risk variables and the build up of the bubble was not a discrete event like the Crisis was. Another potential problem is that the instruments might be correlated with changes NO_x from cars. This would bias second-stage estimates upward if industrial exposure were correlated with car exposure *and* the Crisis also caused a sudden and permanent drop in car usage in the area. The former condition is unlikely given the large area that firms affect, while highways rarely have a significant impact beyond 500 meters (Karner, Eisinger, and Niemeier 2010; Anderson 2015). Furthermore, traffic data show that no significant change in driving patterns coincided with the Crisis.[24]

1.4.2 Using the Chemistry of NO_x and Ozone

Several characteristics of NO_x and ozone make them ideal for identifying how much buyers depend on the visibility of pollutants in their decision making.

First, NO_x and ozone serve as good counterfactual chemicals for one another. They are both lung irritants, but NO_x has a reddish-brown color and noxious smell, while ozone is invisible, has no bad odor, and is far more toxic than NO_x.[25] Thus, if people respond more to NO_x, it is likely because of its greater visibility, while if people respond more to ozone, it is likely because of its greater toxicity.

Second, ozone is the product of NO_x-dependent atmospheric reactions, so the Electricity Crisis exogenously shocked people's exposure to both pollutants. NO_x, a catchall term for NO and NO_2, is emitted directly by combustion processes while ozone is created from NO_x-dependent chemical reactions. These chemical reactions also depend on UV radiation

24. Unreported regressions show traffic patterns had no significant break from trend through the period of the Crisis. I use data from the California Department of Transportation's Freeway Performance Management System (PeMS) for the Bay Area (region 11), 1999–2005. The Bay Area is used because data for Los Angeles only go back to 2001.

25. NO_x and ozone are oxidizing agents, which react with and destroy cells in the lining of the lung, making it more difficult for the lungs to clear foreign particles and bacteria (Chitano et al. 1995). See also Sullivan (1969).

Table 1.1: *Seasonal Trends in Pollution*

	Unconditional Mean			Regression		
	Haze	NO_x	Ozone	Haze	NO_x	Ozone
Q1	0.200	0.359	-0.626			
Q2	-0.616	-0.561	0.701	-0.805***	-0.904**	1.312***
				[0.188]	[0.292]	[0.166]
Q3	-0.329	-0.448	0.467	-0.522**	-0.797**	1.084**
				[0.183]	[0.285]	[0.306]
Q4	0.748	0.654	-0.545	0.552**	0.295**	0.084
				[0.154]	[0.087]	[0.094]

Notes: N=499. Data are monthly averages of hourly readings from the 6 monitors in and near SCAQMD that had readings for all three pollution measures. Each cell is the raw mean of the measure in a quarter or the conditional mean calculated from a regression of the pollution measure on quarter dummies and monitor-year dummies. Sample period is 1991–1997; all monitors have at least 82 of 84 possible monthly observations. Pollution measures have been standardized to have mean 0, standard deviation 1. "Haze" is the coefficient of haze. Standard errors, clustered by monitor, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

from the sun, which varies predictably throughout the year.[26]

Third, the predictable variation in UV radiation across seasons results in different but predictable rates of NO_x-to-ozone conversion throughout the year. Table 1.1 shows these trends using data from pollution monitors on NO_x, ozone, and the coefficient of haze. Columns 1–3 show unconditional means and columns 4–6 show coefficients from regressions of the given pollution metric on dummy variables for each quarter, controlling for monitor-year fixed effects. Standard errors are clustered by monitor. To make comparisons easier, each variable has been standardized to have a mean of 0 and a standard deviation of 1. The patterns in both sets of statistics are the same, regardless of controls: NO_x peaks in Q4 and

26. The relationship of NO_x and ozone can be summarized as two chemical processes:

$$\begin{aligned} hv + NO_2 + O_2 &\rightarrow NO + O_3 \\ NO + O_3 &\rightarrow NO_2 + O_2 \end{aligned}$$

In the first process, an ultraviolet photon (hv), an NO_2 molecule, and free oxygen (O_2) react to form ozone (O_3) and NO. In equilibrium, ozone and NO readily react with each other to reverse this process, leaving no net ozone. Disturbances to this equilibrium, like increasing UV radiation or adding volatile organic compounds (VOCs) that convert NO to NO_2, result in increased ozone levels. See Sillman (1999) for a more detailed discussion of NO_x-ozone reactions.

bottoms out in Q2, with ozone following the opposite pattern.

Because the seasonal variation in NO_x-to-ozone conversion is predictable, it can be used to test whether or not people respond to transitory changes in pollution levels as discussed in Section 1.3.2. This can be done by allowing the effect of aermod_{it} to differ across seasons of the year. Recall that AERMOD uses data on firms' NO_x emissions and is agnostic about how those original emissions change in the atmosphere before arriving at their final destination. This makes aermod_{it} a measure of both NO_x and ozone, depending on the season. In Q2 aermod_{it} will be primarily ozone, while in Q4 it will be primarily NO_x. A larger effect in Q4 would suggest that people respond more to the visible NO_x, while equal coefficients in all quarters would suggest that agents are not affected by salience and may have perfect information.

1.4.3 Estimation Strategy

To measure how prices respond to pollution exposure, I estimate the following equation:

$$\ln p_{it} = \beta \cdot \text{aermod}_{it} + \alpha_i + \delta_t + \sum_k \gamma_{1k} \cdot w_{ik} \cdot t + \sum_k \gamma_{2k} \cdot w_{ik} \cdot t^2 + \varepsilon_{it} \tag{1.8}$$

where p_{it} is the price of house i in quarter t; aermod_{it} is exposure to industrial NO_x-based pollution; α_i are house fixed effects; δ_t are time (quarter-year) effects; $(\gamma_{1k}, \gamma_{2k})$ are coefficients of quadratic time trends for local geographies, defined by a 10-km grid, and local economic conditions that might affect house prices (discussed below); and ε_{it} is the usual residual term.

These controls account for a number of factors that may confound estimates of β, such as amenities not included in the available data and differential trends across local housing markets. The house fixed effects, α_i, capture of all time-invariant characteristics about the house like square footage, number of bedrooms, proximity to the beach, etc. The time effects, δ_t, account for general trends in the housing market over time, as well as seasonal trends within each year (e.g., if houses consistently sell for more during the summer). The local geographic trends allow different parts of the metropolitan area to have different secular trends.[27]

The local trends by economic variables are specifically targeted at concerns related to the housing bubble, which differentially impacted neighborhoods with poor credit. Mian and Sufi (2009) find that zip codes with lower incomes and credit scores were affected more by the expansion of sub-prime credit. If these areas also experienced relatively bigger air quality improvements thanks to the Crisis, the coefficient on aermod_{it} could pick up any increase in house prices due to the expansion of sub-prime credit. To prevent this, I interact the following variables with quadratic time trends: the average loan-to-value ratio for houses sold in the house's census tract in 2000; the average predicted interest rate for mortgages taken out in the house's census tract in 2000; and the median household income in the house's census block group in 2000. The first two variables are averages at the tract level, rather than block group, because they are based on transacted properties in that year, making the smaller block group sample too noisy. The predicted mortgage interest rate data was calculated by DataQuick using proprietary methods and is included in the house data described in Section 1.5.1.

I restrict the analysis to the 11-year period centered on the Crisis: 1995–2005. RECLAIM and the data collection it required was rolled out across firms during 1994, so the first full year of reliable emissions data is 1995. I follow Fowlie, Holland, and Mansur (2012) and set the last year of my analysis to 2005. This avoids the peak and subsequent collapse of the housing bubble. However, because I use the first two years of data in constructing my instruments, the actual regression sample is limited to 1997–2005.

I restrict the region of analysis to the southwest region of SCAQMD territory, roughly between Santa Monica and Huntington Beach (see Figure 1.3a), to minimize measurement error due to geography. All of the major polluters are located in this region and houses

27. Given the large size of the sample region, the ideal geographic unit for these trends would be individual cities, which have economically meaningful boundaries (unlike zip codes) and are generally small but not so small as to be computationally burdensome (unlike tracts and zip codes). Unfortunately, many houses do not have a city listed in the data, and the cities of Los Angeles and Long Beach cover a large portion of the sample region while also having a great deal of within-city heterogeneity. To overcome these issues, I use a 10-km grid, shown in Figure A.1, which is aligned to preserve as many city boundaries as possible. This grid results in 17 different areas that each get their own quadratic time trend.

farther away from the pollution sources are likely to have less actual exposure from the firms and more noise in the modeling prediction, decreasing the signal-to-noise ratio of the pollution measure. Predicting the pollution distribution is also more complicated farther inland because of the San Gabriel and Santa Ana Mountains, which can act like a dam, collecting pollution blown from the coasts. To avoid these problems, I restrict my sample to houses within 10 kilometers of a major electric firm in Los Angeles or Orange County.[28]

Because of the previously discussed concern about omitted variables, I estimate Equation (1.8) using two-stage least squares (2SLS) and limited-information maximum likelihood (LIML). As discussed in Section 1.4.1, the main set of instruments is $\{\text{aermod\_pre}_i \times \delta_y\}_{y=1998}^{2005}$, which I refer to as the "annual" instruments and which have the following first stage:

$$\text{aermod}_{it} = \sum_{y=1998}^{2005} \left(\text{aermod\_pre}_i \times \delta_y\right) \cdot \pi_y + \nu_i + \nu_t + \sum_k \eta_{1k} \cdot w_{ik} \cdot t + \sum_k \eta_{2k} \cdot w_{ik} \cdot t^2 + \mu_{it} \quad (1.9)$$

The estimates of π_y can be plotted to verify the common trends assumption that the instruments only affect aermod_{it} and $\ln p_{it}$ through the Crisis. I also use an alternative "post" instrument, $\text{aermod\_pre}_i \times \text{post}_t$ where $\text{post}_t = \mathbf{1}\{t \geq 2001\}$, instead, which treats the Crisis as a difference-in-differences with variable treatment intensity. This provides a convenient summary of the average effect of the Crisis on exposure and prices.

To test for pollutant visibility effects, I allow the effect of aermod_{it} in Equation (1.8) to vary by quarter of year, as discussed in Section 1.4.2:

$$\ln p_{it} = \sum_q \beta_q \cdot (\text{aermod}_{it} \times Q_q) + \alpha_i + \delta_t + \sum_k \gamma_{1k} \cdot w_{ik} \cdot t + \sum_k \gamma_{2k} \cdot w_{ik} \cdot t^2 + \varepsilon_{it} \quad (1.10)$$

where Q_q is an indicator equal to 1 if t is the q-th quarter of the year.

28. I also include in this group the southwestern most firm in the area in order to include the Palos Verdes Peninsula in the regression sample (see Figure A.1).

1.5 Data

1.5.1 Houses

Data on home sales and housing characteristics come from county registrar and assessors' offices and were collected by DataQuick, Inc. The data include any property that has been assessed and most sales, refinances, and foreclosures in California after 1990. Data for each property includes square footage, lot size, number of bedrooms and bathrooms, and the year the property was built. Each sale or refinance includes the value of the mortgage and any additional loans taken against the property, as well as interest rates as estimated by DataQuick using proprietary methods. Latitude and longitude are also included for each property.

I drop sales that fall outside normal market transactions and which may not accurately reflect the market's valuation of the house. Specifically, all transactions must be arms-length, non-distressed sales (i.e., no foreclosure sales or short sales) with a price of at least $10,000. I also drop a sale if the property transacted within the previous 90 days, as many of these transactions are duplicates. The sample is also restricted to homes built before 1995 to preclude direct sales from developers to consumers. The top 0.1% of sales are winsorized.

Table 1.2 shows average sale price, house hedonics, and quarter of sale broken down by the number of times a house transacted during the sample period. House prices are deflated to real 2014 dollars using the all-items CPI. Houses are are not used in summary statistics or regressions if they fall outside the sample region described in Section 1.4.3.

1.5.2 Firms

There are several components to the firm-level data, which cover firm emissions over time, the firm's name and location, and physical characteristics of the firm's polluting equipment. The firm data also include information about RECLAIM Trading Credit (RTC) allocations and subsequent trades.

Most of the data come from SCAQMD via public records requests (SCAQMD 2015a).

Table 1.2: *House Summary Statistics*

	Never Sold	Sold Once		Repeat Sales	
		Pre	Post	Pre	Post
Sale Price		394,621	541,016	420,397	603,089
		(284,495)	(357,224)	(303,028)	(395,970)
Lot Size	19,963	14,831	19,454	19,444	14,650
	(943,394)	(812,098)	(918,742)	(992,280)	(807,084)
Square Feet	1,537	1,611	1,534	1,573	1,491
	(647)	(721)	(689)	(707)	(654)
Year Built	1950	1952	1950	1951	1950
	(15.15)	(15.61)	(15.77)	(16.96)	(16.78)
Bedrooms					
1	0.01	0.01	0.01	0.01	0.02
2	0.23	0.22	0.24	0.25	0.27
3	0.48	0.48	0.49	0.49	0.49
4	0.22	0.23	0.21	0.21	0.19
5+	0.05	0.05	0.05	0.04	0.03
Bathrooms					
1	0.34	0.29	0.33	0.31	0.35
2	0.47	0.47	0.46	0.45	0.45
3	0.13	0.16	0.13	0.15	0.13
4+	0.03	0.04	0.04	0.05	0.04
Sold in Quarter					
1		0.19	0.22	0.20	0.21
2		0.28	0.27	0.29	0.28
3		0.28	0.28	0.28	0.27
4		0.25	0.24	0.24	0.23
Times Sold				2.14	
				(0.38)	
Total Properties	240,110	84,011		19,539	

Notes: Summary statistics from regression sample as described in Section 1.5.1. Table lists sample means with standard deviations given in parentheses.

These data include each firm's name, address, SCAQMD-assigned ID number, the mass of NO_x the firm emitted every quarter from 1994 to 2014, and all relevant RTC data, including initial allocation of RTCs, the quantity, price, and vintage of exchanged RTCs, and the ID numbers of participating firms. Firms' operating addresses were geocoded to get latitude and longitude to represent the location of the firm's smoke stacks, which are required by AERMOD and other location-based calculations (see Section A.1.1 for more details).

AERMOD requires data on the physical characteristics of firms' polluting equipment (smoke stack height and diameter, and temperature and velocity of gas exiting the smoke stack), which I take from the National Emissions Inventory (NEI). Regulators often collect these data specifically to run atmospheric dispersion models like AERMOD, but the data collected by SCAQMD could not be made available (SCAQMD 2015b). However, the National Emissions Inventory (NEI) has these data for many firms along with each firm's name, address, SIC, and the type of combustion process behind each stack. I matched most firms to the NEI by reconstructing the NEI-specific ID number from the SCAQMD ID number and other administrative variables, and I validated these matches using fuzzy string matching on firm names and addresses. Any remaining firms were matched via fuzzy string matching and manually checked. For firms with missing stack data, I impute values using the firm's SIC and the stack's equipment-type code (SCC). Details of the imputation process and the construction of the firm-level data in general are outlined in Section A.1.

Table 1.3 gives summary statistics by industry (4-digit SIC) on emissions, smoke stack parameters, electric-generator status, average distance to the nearest meteorological station, and the number of firms in each industry group.

1.5.3 Meteorology and Pollution Monitor Data

Data on local meteorological conditions come from SCAQMD. Before building new polluting equipment, firms must submit an impact report to SCAQMD using AERMOD to show how the new equipment will impact ambient pollution levels. To facilitate the making of these

Table 1.3: *Firm Summary Statistics by Industry*

(a) *Tons of NO_x Emitted*

Industry	Mean		Median		Share of Total	
	1998	2002	1998	2002	1998	2002
Petroleum Refining	665.20	479.08	818.57	492.87	51.1%	63.4%
Electric Services	213.19	60.14	100.73	48.94	22.9%	11.1%
Glass Containers	199.27	107.82	145.04	77.92	4.6%	4.3%
Crude Petroleum and Natrual Gas	36.43	8.86	5.72	1.42	3.6%	1.5%
Other Petroleum and Coal Products	321.18	301.88	321.18	301.88	2.5%	4.0%
Steam and Air-Conditioning Supply	38.80	5.65	14.48	3.71	1.8%	0.4%
Other Industrial Inorganic Chemicals	39.70	37.01	34.50	43.59	1.2%	1.5%
Secondary Smelting and Refining	50.84	27.34	52.62	27.63	1.2%	1.1%
Flat Glass	116.21	50.52	116.21	50.52	0.9%	0.7%
Gas and other Services	107.87	9.45	107.87	9.45	0.8%	0.1%
Other Industries	9.61	6.70	4.64	2.91	9.3%	11.9%
All firms	71.47	39.99	6.98	4.26	100.0%	100.0%

(b) *Physical Characteristics*

Industry	Smoke Stack				Dist. to	
	Height (m)	Diameter (m)	Velocity (m/s)	Gas Temp. (K)	Met. Site (km)	Firms
Petroleum Refining	30.10	1.59	11.83	548.83	6.52	10
Electric Services	40.84	3.69	19.49	481.12	7.46	14
Glass Containers	26.09	1.23	13.41	495.60	7.82	3
Crude Petroleum and Natrual Gas	6.85	0.34	13.56	595.11	6.09	13
Other Petroleum and Coal Products	14.66	0.63	14.56	341.32	6.06	1
Steam and Air-Conditioning Supply	19.68	0.83	12.67	468.18	6.76	6
Other Industrial Inorganic Chemicals	35.23	0.97	11.87	540.11	6.09	4
Secondary Smelting and Refining	9.35	0.69	14.11	406.11	5.52	3
Flat Glass	10.97	1.28	13.60	547.04	5.34	1
Gas and other Services	18.29	4.36	22.49	727.59	6.25	1
Other Industries	12.43	0.82	10.01	486.12	6.17	126
All firms	16.14	1.08	11.46	497.27	6.31	182

reports, SCAQMD makes AERMOD-ready meteorological data available on its website.[29]

These data were gathered by 27 meteorological stations throughout SCAQMD.[30] The data consist of hourly observations for temperature, mean and standard deviations of wind speed and direction at multiple altitudes, and other variables described in Section 1.2.2. Each station provides at least three years of data between 2006 and 2012. While these stations were not in operation at the time of the Crisis, wind patterns at the given locations are very stable over time.

Data from air pollution monitoring stations comes from the California Air Resources Board (CARB). They include hourly measures of ozone, NO_x, and the coefficient of haze (COH), which is a measure of visibility interference in the atmosphere. I aggregate the hourly measures to daily and then monthly averages following Schlenker and Walker (Forthcoming).

1.5.4 AERMOD-based Measure of Exposure

I use AERMOD, which maps firm-level output to house-level exposure, to construct a measure of a house's exposure from all industrial sources. Software for using AERMOD is available on the EPA's website and includes documentation, Fortran source code, and pre-compiled executables for Windows.[31]

As discussed in Section 1.2.2, house i's exposure to NO_x emissions from firm f at time t can be written $\text{NOx}_{ft} \cdot h(d_{fi}, \theta_{fi}; \boldsymbol{S}_f)$, where $\boldsymbol{S}_f$ contains information on the firm's smoke stacks, as well as local meteorological conditions. The data I use for NOx_{ft} and $\boldsymbol{S}_f$ are described in Sections 1.5.2 and 1.5.3. A firm's meteorological data is taken from the nearest meteorology monitor. The values for (d_{fi}, θ_{fi}) are calculated by AERMOD from firms and

29. The data are most easily accessible via the SCAQMD website: http://www.aqmd.gov/home/library/air-quality-data-studies/meteorological-data/data-for-aermod

30. The location of these stations is mapped in Figure A.2.

31. See http://www.epa.gov/scram001/dispersion_prefrec.htm. I use AERMOD version 13350, compiled using Intel Fortran Compiler 15.0 for Linux and run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University.

houses' latitude and longitude. AERMOD then outputs aermod_{ift}, the house's exposure to the firm's emissions. The house's total exposure to industrial NO_x emissions is simply $\text{aermod}_{it} = \sum_f \text{aermod}_{ift}$.

For block group–level exposure, I first calculate exposure at the block level, then calculate the population-weighted average for each block group. At the block level, I use the process described above for houses, substituting house-specific latitude and longitude for the Census-provided internal point for each block.[32] This is a more attractive approach than using the block group's internal point because it accounts for heterogeneity in population and exposure across the block group and is a closer approximation to the average exposure to the block group's residents.

Because AERMOD loops over all firms, houses, and meteorological data, it is very computationally intensive for such a large sample, so I impose several restrictions on the data to make calculation more feasible.[33] First, I only calculate a firm's exposure to houses that are within 20 kilometers of the firm and set exposure outside this radius to zero. Second, I use one year of meteorological data, 2009, which is also the only year during which all of the meteorological stations described in Section 1.5.3 were operating. Third, I construct an arbitrary 100-meter grid by rounding each house's UTM coordinates to the nearest 100 meters and calculate the exposure value at the center of each grid square. Houses are then assigned exposure according to the grid square they occupy.

---

32. Analyses using Census geographies like block groups or ZCTA's often use the "centroid" of the geography as its the representative point in space. However, the Census Bureau is very particular to note that because these geographies are not convex, the true centroid may lie outside the geography of interest. As a solution, the Census Bureau calculates "internal points," which are constrained to be inside the geography.

33. Even with these restrictions, the model takes approximately 210 CPU days to process all the data.

Table 1.4: *Pollution's effect on House Price, OLS*

	(1)	(2)	(3)	(4)
Aermod	-0.0507***	-0.0139***	-0.0030***	-0.0033***
	[0.0019]	[0.0009]	[0.0006]	[0.0010]
Controls	N	Y	Y	Y
Fixed Effects	None	None	BG	House
R^2	0.0461	0.7678	0.8649	0.9483
N	118,522	118,522	118,522	41,771

Notes: Outcome variable is ln house price. Controls include year-quarter effects, quadratic time trends by local geography and year 2000 SES variables, and hedonics: lot size, bedrooms, bathrooms, square feet. Observations absorbed by fixed effects are dropped. Standard errors, clustered at 100-m grid, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

1.6 Results

1.6.1 OLS Estimates of Prices and Exposure

I start with simple OLS regressions of log house price on pollution exposure. All regressions are clustered at the 100-m grid used to calculate aermod_{it} (see Section 1.5.4).

Column 1 of Table 1.4 shows the naïve univariate regression of prices on exposure. Column 2 adds the year-quarter effects, geographic time trends, and SES time trends from Equation (1.8), as well as explicit hedonic controls: number of bedrooms, number of bathrooms, square footage, and lot size. Adding time trends and hedonics in Column 2 reduces the effect of aermod on prices by about 75%, suggesting that the coefficient in Column 1 is picking up the fact that houses in polluted areas have different characteristics (e.g., are smaller), and that secular trends affect both pollution and house prices. Column 3 adds block group fixed effects which further reduces the effect of exposure, suggesting that neighborhood-level characteristics are also important. Column 4, the preferred specification defined in Equation (1.8), trades the block-group effects and hedonics for property-specific fixed effects. The estimate is not dramatically different from Column 3, though it is slightly larger.

In aggregate, the results of Table 1.4 suggest that omitted variables are a potentially serious problem when measuring the effect of pollution exposure.

1.6.2 Event Studies around the Crisis

For an instrument based on the Crisis to yield consistent estimates of the effect of pollution exposure, the instrument should have no effect before the Crisis began. This common trends assumption requires that, had the exogenous shock not taken place, individuals of varying treatment intensity would have maintained their status quo. The common trends assumption can be assessed by plotting the event study coefficients from Equation (1.9), the first stage, and the analogous coefficients from the regression of price on the instruments. For the first stage, each coefficient $\hat{\pi}_y$ can be interpreted as the change in relative exposure across areas with different initial exposure levels, aermod_pre_i, relative to the difference across these areas in the omitted year, 2000. For example, if relative exposure does not change between 2000 and 2001, $\hat{\pi}_{2001} = \hat{\pi}_{2000} = 0$, since π_{2000} is omitted and thus constrained to be zero. If, on the other hand, relative exposure decreases in high aermod_pre_i areas, $\hat{\pi}_{2001}$ will be negative ($\hat{\pi}_{2001} < \hat{\pi}_{2000} = 0$).

As Figure 1.6 shows, it appears that the common trends assumption holds and the Crisis had a large effect on both exposure and prices. For both exposure and prices, Figure 1.6 shows that aermod_pre_i had no effect before the Crisis, with sharp effects afterwards, suggesting that the Crisis makes a good natural experiment and that house prices respond sharply to exposure levels. As Figure 1.4 and Figure 1.5 from Section 1.4.1 show, the Crisis hit most firms in mid- to late-2000. In Figure 1.6, we see no significant change in exposure or price between 1997 and 2000. In 2001, relative exposure suddenly drops for high aermod_pre_i properties and continues to decline slightly afterward, consistent with firms' drop in emissions. Similarly, house prices jump at the same time exposure falls and, in a noisy mirror image of the exposure trend, continue to appreciate slightly over time.

The timing of these sharp jumps immediately after the Crisis also suggests that the Crisis, and not coincidental secular trends, is driving these changes.

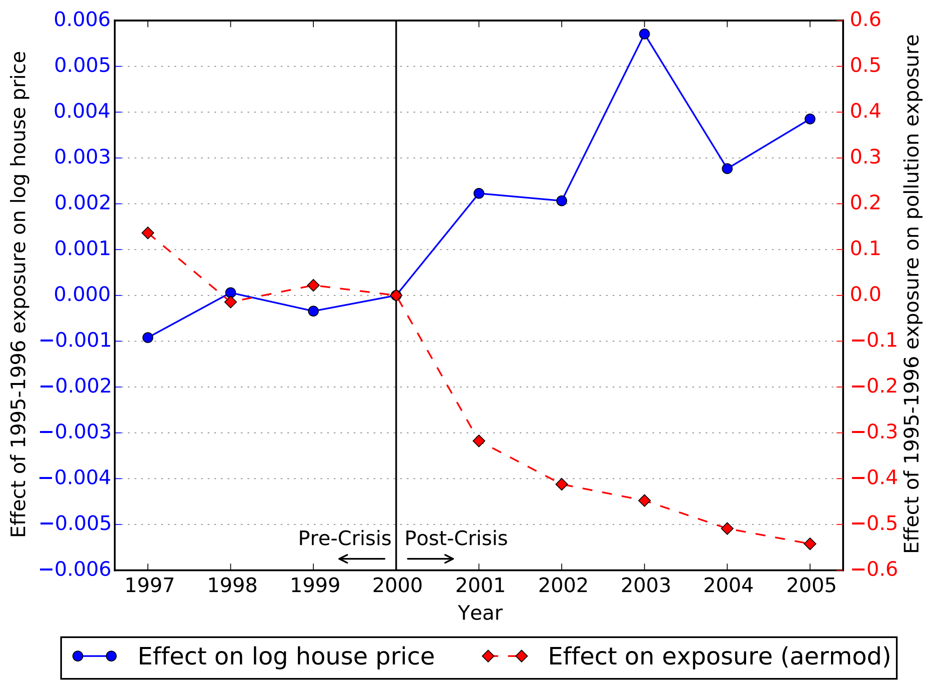

Figure 1.6: *Pollution Exposure and House Prices around the Crisis*

Notes: Plotted points are coefficients from a regression of the specified outcome on the "annual" instruments: aermod_pre interacted with year dummies. See Equation (1.9). Here, the year of the Crisis, 2000, is the omitted group. Sample and other controls as in Table 1.5, columns 2–6. aermod_pre is the average of the aermod exposure variable for 1995 and 1996. Average value of aermod_pre is 5.172.

Table 1.5: *Pollution's effect on House Price, Instrumental Variables*

	(1) ln Price	(2) ln Price	(3) Aermod	(4) ln Price	(5) ln Price	(6) ln Price
Aermod				-0.0073***	-0.0073***	-0.0073***
				[0.0024]	[0.0023]	[0.0024]
Aermod_pre×post	0.0033***	0.0032***	-0.4328***			
	[0.0005]	[0.0008]	[0.0748]			
Aermod_pre	-0.0029**					
	[0.0012]					
Fixed Effects	BG	House	House	House	House	House
Method	OLS	OLS	OLS	2SLS	2SLS	LIML
IV set				Post	Annual	Annual
κ				1	1	1.0003
1st Stage F-stat				6388	932	932
R^2	0.865	0.948	0.911			
N	118,522	41,771	41,771	41,771	41,771	41,771

Notes: Sample average of aermod_pre is 5.172. In addition to fixed effects, controls include year-quarter effects and quadratic time trends by local geography and year 2000 SES variables (see Section 1.4.3). For full output of columns 2–4, see Table A.1. Column 1 also includes the following hedonic controls: lot size, bedrooms, bathrooms, square feet. "Post" IV is aermod_pre × post, "Annual" IV is aermod_pre interacted with year dummies. First-stage F stat assumes homoskedasticity. Observations absorbed by fixed effects are dropped. Standard errors, clustered at 100-m grid, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

1.6.3 Instrumental Variables Estimates of Price Effects

Columns 1–3 of Table 1.5 present first-stage and reduced-form estimates using the "post" instrument, which provides a more concise summary of the effects of the Crisis. In order to show the main effect of aermod_pre$_i$ on prices, Column 1 uses the same hedonics and block group fixed effects as Table 1.4, Column 3. The main effect of aermod_pre is -0.0029, suggesting that, on average, houses with 1 additional μg/m^3 of pre-Crisis exposure were valued 0.29% lower than comparable houses. The effect of the "post" instrument, aermod_pre × post, is 0.0033, suggesting that the value of previously high-pollution houses saw their value equalize with houses that had low pollution throughout the period.

Columns 2 and 3 show the estimated effect of the post instrument on prices and exposure, respectively, using the preferred specification with property fixed effects (see Table A.1

for full output). The reduced form estimate in Column 2 is similar to that of Column 1 and shows that house prices increased 0.3% per unit of treatment intensity. At the average value of aermod_pre, this translates to a 1.7% increase in value, or $7,324 for a house of average value in 2000. Similarly, Column 3 shows that exposure decreased by 0.433 μg/m^3 NO_x/ozone for every unit of treatment intensity, or 2.24 μg/m^3 for the average value of aermod_pre$_i$.

Columns 4 and 5 of Table 1.5 present the 2SLS results using the "post" instrument, aermod_pre$_i \times$ post$_t$, and the annual set of instruments, aermod_pre$_i \times \delta_y$, respectively. (The full results of Column 4 are shown in Table A.1.) These results are almost identical and suggest that an additional μg/m^3 of exposure to NO_x emissions decreases the value of a house by about 0.7%. This translates to a MWTP to reduce exposure of $3,272 per μg/m^3. While not directly comparable, this figure more than covers the $PM_{2.5}$ mortality cost of $1,600 per μg/m^3 borne by a family of two adults and one child discussed in Section 1.2.

The 2SLS estimates do not appear to suffer from weak-instruments bias, as evidenced by the instruments' partial F statistics from the first stage and the LIML results also presented in Table 1.5. Following Stock and Yogo (2002) and Stock, Wright, and Yogo (2002), I use the instruments' partial F statistic in the first stage to assess whether the instruments are weak. The F statistics, assuming spherical errors, for the post and annual instruments are 6,323 and 923, respectively, leaving little worry about a weak instruments problem.[34] The LIML estimates in column 3 provide further evidence against weak instruments because the LIML estimator is median-unbiased and thus more reliable than 2SLS when instruments are weak (Stock, Wright, and Yogo 2002). If the LIML estimates differ from 2SLS, concerns about weak instruments could be warranted. However, that does not appear to be the case here as the 2SLS and LIML estimates are virtually identical.

---

34. Following Stock and Yogo (2002) and Stock, Wright, and Yogo (2002), it has become standard practice to measure the strength of excluded instruments using the partial F statistic from the first stage. However, the usual rules of thumb from Stock, Wright, and Yogo assume spherical error terms. The correct test statistic for robust first-stage F stats is an open topic of research (see, e.g., Montiel Olea and Pflueger 2013). Therefore, I follow the approach of Coglianese et al. (2015) and report the non-robust F statistics in Table 1.5 for comparison against the usual rule of thumb.

Table 1.6: *Robustness to Spatially Correlated Error Terms*

	Std. Err.	p-value
Baseline (clustered)	0.0024	0.0022
SHAC by Bandwidth (m)		
200	0.0025	0.0036
400	0.0028	0.0086
600	0.0031	0.0178
800	0.0033	0.0279
1000	0.0035	0.0362
1200	0.0036	0.0414
1400	0.0037	0.0452
1600	0.0037	0.0480

Notes: N=41,771. Each row re-estimates the standard error of aermod in the 2SLS regression in Table 1.5, column 4 using the non-parametric Spatial HAC (SHAC) method of Conley (1999) and Kelejian and Prucha (2007). Kernel used is a triangle with the listed bandwidth. Clustered standard error from baseline regression is given on the first row.

The 2SLS estimates are also robust to arbitrary spatial correlation across the error terms. This is shown in Table 1.6 by re-estimating the standard errors from the preferred specification (Table 1.5, column 1) using the spatial HAC (SHAC) variance-covariance estimator of Conley (1999) and Kelejian and Prucha (2007).[35] I use a triangle kernel with bandwidths from 200 meters to 1600 meters (1 mile) and list the standard error and corresponding p-value for each bandwidth. The p-value at each bandwidth is less than 0.05, suggesting that the estimates are indeed statistically significant. The standard errors also increase with bandwidth at a decreasing rate, further suggesting that the estimates are credibly precise.

These above estimates are corroborated by the results of Chapter 2, which estimates the

35. SHAC standard errors can be thought of as an extension of Newey–West standard errors from discrete time to continuous distance. Specifically,

$$\mathrm{Var}(\hat{\beta}) = \left(\sum_i x_i x_i'\right)^{-1} \left(\sum_i \sum_j K(d_{ij}) x_i \hat{\varepsilon}_i \hat{\varepsilon}_j x_j'\right) \left(\sum_i x_i x_i'\right)^{-1}$$

where K is some kernel and d_{ij} is some metric of the distance between units i and j.

effects of the Crisis on block group–level rents and demographics. The estimated reduced-form effect on rents is 0.0031, very similar to the estimate here for house prices. Additionally, the paper finds evidence of a large sorting response following the Crisis, confirming that the change in prices is due to an actual change in amenities rather than contemporaneous secular trends in the housing market.

1.6.4 Comparison to Standard Methods

To verify that the large MWTP estimate found in Section 1.6.3 is not being driven by a peculiarity of the data or natural experiment, I re-estimate MWTP using non–wind-based instruments standard in the literature. I use two standard ways of constructing an instrument based on the Crisis: geographic difference-in-differences and kernel-based measures of exposure similar to those used by Banzhaf and Walsh (2008).

Geographic Diff-in-diff and Interpolation

The first standard research design is the geographic diff-in-diff. The equation to estimate is similar to Equation (1.8), but each pair of house i and firm f is treated as a separate observation so that the same sale price p_{it} can appear with multiple firms:

$$\ln p_{ift} = \text{near}_{if} \times \text{post}_t \cdot \beta + \alpha_{if} + \mathbf{X}_{it}\mathbf{\Gamma} + \varepsilon_{ift} \tag{1.11}$$

where the entity fixed effects are now house-firm effects instead of house effects; $\mathbf{X}_{it}$ includes the same time and demographic controls as Equation (1.8); and near_{if} is a dummy variable for whether house i is within the set treatment radius of firm f.[36] I estimate this model on the full study sample twice, once with a 1-mile treatment radius and a 2-mile control, and once with a 2-mile treatment and 4-mile control.

The reduced-form estimates, shown in columns 1 and 4 of Table 1.7, are small, imprecise, and have different signs. For the 1-mile treatment, the average effect of the Crisis on log price is 0.0049, about one third the size of the effect estimated in Table 1.5, column 2, for a

36. For a similar application, see Currie et al. (2015).

Table 1.7: *Price Effects with Geographic Diff-in-diff*

	(1)	(2)	(3)	(4)	(5)	(6)
	0–1 vs. 1–2 miles			0–2 vs. 2–4 miles		
	ln Price	Aermod	ln Price	ln Price	Aermod	ln Price
Near×post	0.0049	-0.4991***		-0.0011	0.0237	
	[0.0049]	[0.0566]		[0.0023]	[0.0222]	
Aermod			-0.0098			-0.0461
			[0.0098]			[0.1029]
Method	OLS	OLS	2SLS	OLS	OLS	2SLS
R^2	0.9453	0.9095		0.9416	0.9104	
N	92,807	92,807	92,807	430,836	430,836	430,836

Notes: Unit of observation is house-firm-quarter. Near=1 for houses closer to firm, e.g., 0–x miles as specified. Controls include house-firm effects, year-quarter effects, and local quadratic time trends. Observations absorbed by fixed effects are dropped. Standard errors, clustered at 100-m grid, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

house of average treatment intensity (0.017). The 2-mile estimate implies that houses close to firms lost value because of the Crisis, but this estimate is also imprecise.

The derivation of geographic diff-in-diff bias in Section 1.2.1 predicts that the first-stage and reduced-form estimates should have the same bias and that, with a good measure of exposure, the second-stage estimate should be unbiased, though potentially noisy. To test this, I use the firm-specific exposure measure aermod_{ift} as the endogenous regressor. For the 1-mile treatment radius, the biases appear to be roughly equal. The reduced-form effect is 32% of the average reduced-form effect found in Table 1.5, column 2, while the first stage effect is 22% of its AERMOD-based equivalent. Consequently, the second stage coefficient, -0.0098, is similar to the estimates in Table 1.5 but very imprecise. For the 2-mile treatment, the reduced-form and first-stage estimates recover only 7% and 1% of the wind-based IV estimates, respectively, and all three estimates are imprecise.

For a more direct comparison with prior literature, I also estimate geographic diff-in-diffs using interpolated NO_x and ozone from pollution monitors and present the results in Table 1.8. As before, the interpolation is calculated using inverse distance weighting using

Table 1.8: *Price Effects with Geographic Diff-in-diff and Interpolation*

	A. 1-mile treatment, 2-mile control						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ln Price	NO_x	ln Price	ln Price	Ozone	ln Price	ln Price
Near×post	0.0056	0.3398			-0.0846		
	[0.0082]	[0.4708]			[0.1125]		
NO_x			0.0164	-0.0070			
			[0.0328]	[0.0065]			
Ozone						-0.0658	0.0028
						[0.1308]	[0.0225]
Method	OLS	OLS	2SLS	2SLS	OLS	2SLS	2SLS
IV Set			Post	Annual		Post	Annual
1st Stage F-stat			0.9	2.0		0.9	3.2
	B. 2-mile treatment, 4-mile control						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ln Price	NO_x	ln Price	ln Price	Ozone	ln Price	ln Price
Near×post	-0.0083**	-1.0768***			0.2228***		
	[0.0034]	[0.1860]			[0.0452]		
NO_x			0.0077**	0.0018			
			[0.0034]	[0.0023]			
Ozone						-0.0373**	0.0037
						[0.0175]	[0.0051]
Method	OLS	OLS	2SLS	2SLS	OLS	2SLS	2SLS
IV Set			Post	Annual		Post	Annual
1st Stage F-stat			50.9	12.1		38.5	39.8
	C. 3-mile treatment, 6-mile control						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ln Price	NO_x	ln Price	ln Price	Ozone	ln Price	ln Price
Near×post	-0.0017	-0.5263***			0.1365***		
	[0.0022]	[0.1082]			[0.0276]		
NO_x			0.0033	0.0051			
			[0.0043]	[0.0037]			
Ozone						-0.0126	0.0018
						[0.0166]	[0.0088]
Method	OLS	OLS	2SLS	2SLS	OLS	2SLS	2SLS
IV Set			Post	Annual		Post	Annual
1st Stage F-stat			26.4	4.8		32.1	15.0

Notes: N for each subtable is 50,746; 264,234; and 423,945, respectively. Unit of observation is house-firm-quarter. NO_x and ozone exposure interpolated from monitors using inverse distance weighting. Near=1 for houses within specified treatment radius. Sample restricted to houses within specified control radius. IV Set "Post" is Near×post. IV Set "Annual" is Near times year dummies. 1st Stage F-stat assumes spherical errors. Controls include house-firm effects, year-quarter effects, and quadratic time trends by local geography and year 2000 SES variables. Observations absorbed by fixed effects are dropped. Standard errors, clustered at 100-m grid, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

monitors with full NO_x and ozone coverage during the sample period that are no more than 10 km from the point being interpolated. Estimates are very sensitive to the treatment and control radii and the instruments used and are generally imprecise or have the wrong sign. The second-stage estimate of ozone's effect on prices in sub-table B, column 6 is the only second-stage estimate that has the correct sign and is precisely estimated. However, it is not robust to the choice of instruments, and the estimate using the "annual" instrument in column 7 is imprecise with the wrong sign.

Kernel-based Exposure

The second non–wind-based research design uses radial kernel densities to map firm emissions to local exposure. Specifically, I use a triangle kernel with 5-km bandwidth and a uniform kernel with 2-km bandwidth as the proxy for the spatial distribution h instead of AERMOD. This is similar to the approach taken by Banzhaf and Walsh (2008), who use the equivalent of a uniform kernel with a 1 mile (1600 meter) bandwidth. The kernel approach should be an improvement over the geographic diff-in-diff because it can account for neighboring firms' overlapping treatment areas. To make the unit-less kernel-based variables comparable to the AERMOD measure, I re-scale them so that their sample average is the same magnitude as the sample average of aermod_{it}. Once again, the estimation equation is almost identical to Equation (1.8), except that the exposure measure and instruments are constructed using the relevant kernel density instead of AERMOD. Table 1.9 presents the results, with the triangle-based regressions in sub-table A and the uniform-based regressions in sub-table B.

The kernel-based estimates, shown in Table 1.9, are also small and imprecise. Column 1 of each sub-table shows the reduced form estimates, which are small and imprecise, with the triangle-based estimate having the wrong sign. Column 2 shows the first stage using aermod_{it} as the endogenous regressor, which is included to be more comparable to my preferred specification and to overcome the fact that the kernel variables have an arbitrary scale. In all cases the excluded instruments are defined using the kernel-based exposure.

Table 1.9: *Price Effects with Kernel-defined Instruments and Exposure*

	A. Triangle Kernel (5-km band)				
	(1)	(2)	(3)	(4)	(5)
	ln Price	Aermod	Triangle	ln Price	ln Price
Triangle_pre× post	-0.0002	-0.1005***	-0.3830***		
	[0.0007]	[0.0090]	[0.0113]		
Aermod				0.0021	
				[0.0071]	
Triangle					0.0006
					[0.0019]
Method	OLS	OLS	OLS	2SLS	2SLS
R^2	0.948	0.888	0.932		
	B. Uniform Kernel (2-km band)				
	(1)	(2)	(3)	(4)	(5)
	ln Price	Aermod	Uniform	ln Price	ln Price
Uniform_pre× post	0.0001	-0.0479***	-0.4071***		
	[0.0003]	[0.0064]	[0.0212]		
Aermod				-0.0026	
				[0.0071]	
Uniform					-0.0003
					[0.0008]
Method	OLS	OLS	OLS	2SLS	2SLS
R^2	0.948	0.888	0.906		

Notes: N=41,771. Sample average of triangle_pre is 2.303. Sample average of uniform_pre is 1.683. Controls include house effects, year-quarter effects, and local quadratic time trends. Observations absorbed by fixed effects are dropped. Standard errors, clustered at 100-m grid, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

These estimates are imprecise and again imply a much smaller average effect than Table 1.5, with neither effect being more than 10% of the wind-based result. Column 3 shows the first-stage regressions using the kernel-based exposure measure, which are precise but hard to compare to Table 1.5 because of the scaling issue. Columns 4 and 5 show the 2SLS estimates using aermod_{it} and kernel-defined exposure, respectively, as the endogenous regressors. When using instruments based on the triangle kernel, the estimates have the wrong sign due to the wrong-signed first stage. When using uniform-based instruments, the estimate in column 4 is almost 50% of the preferred AERMOD-based estimates in Table 1.5, but is imprecise, while the estimate in column 5 is both economically and statistically insignificant.

Summary and Comparison to Prior Research

Table 1.10 summarizes all the estimates from above along with previously discussed estimates from the literature. The first column lists the model or paper that generated the estimate; the second column lists the estimated effect of the Crisis on average house prices for models from this paper; and the third column lists the estimated MWTP for a 1 $\mu g/m^3$ reduction in pollution. For the models estimated in this paper, the pollutant is NO_x and/or ozone, while for Smith and Huang (1995) and Chay and Greenstone (2005) it is TSP, and for Bayer, Keohane, and Timmins (2009) it is PM_{10} (see footnote 4). For this comparison, I do not combine non–wind-based designs with aermod_{it} in any way, as the point of the comparison is to gauge the importance of the wind. Hence, there are no MWTP estimates from the geographic diff-in-diff models because the geographic diff-in-diff has no independent measure of exposure. I also do not include the interpolated regressions from Table 1.8 because the are based on a slightly different geographic sample.

There are several points of interest in Table 1.10 that support the predictions made in Section 1.2.1 that standard estimates may be biased downward. First and foremost, the AERMOD-based estimates dwarf all other estimates in magnitude and precision. Second, the uniform kernel estimate, though imprecise, is not dissimilar from prior research. Third,

Table 1.10: *Comparison of Pollution Estimates Across Models*

	Model/Paper	Crisis' Effect on Avg. Price	MWTP
		Standard models	
(1)	Geo DD (1 mile)	\$1,438	
(2)	Geo DD (2 miles)	−\$589	
(3)	Triangle kernel	−\$217	−\$246
(4)	Uniform kernel	\$95	\$138
		Prior Research	
(5)	SH 1995 (3rd q-tile)		\$233**
(6)	SH 1995 (mean)		\$260**
(7)	CG 2005		\$191**
(8)	BKT 2009		\$130***
(9)	BKT 2009 (w/ moving)		\$350**
		Wind-based model	
(10)	Aermod	\$7,324***	\$3,272***

Notes: Each row is taken from a different research design. "Effect of Crisis" is the reduced form effect of the Electricity Crisis calculated at sample averages. For estimates from other papers, the authors' stated preferred estimate is used. Geo DD, Triangle, and Uniform rows use only results specific to those research designs, i.e., no first or second stage using Aermod-based exposure. Significance levels taken from original sources: ** $p < .05$, *** $p < .01$
Row 1: Table 1.7, col 1
Row 2: Table 1.7, col 4
Row 3: Table 1.9A, cols 1 & 5
Row 4: Table 1.9B, cols 1 & 5
Row 5: Smith and Huang (1995), abstract, meta-analysis
Row 6: Smith and Huang (1995), abstract, meta-analysis
Row 7: Chay and Greenstone (2005), Table 5A, col 4
Row 8: Bayer, Keohane, and Timmins (2009), Table 6, col 2
Row 9: Bayer, Keohane, and Timmins (2009), Table 6, col 4; accounts for moving costs
Row 10: Table 1.5, cols 2 & 4

the instrumental variables estimates from prior research (Chay and Greenstone 2005; Bayer, Keohane, and Timmins 2009) are not dramatically different from the prior OLS estimates (Smith and Huang 1995)—the OLS estimates fall between the IV estimates. All of these observations are consistent argument in Section 1.2 that standard methods of measuring exposure are biased, even when quasi-experiments and instrumental variables are used.

1.6.5 Evidence of Visibility Effects

As discussed in Section 1.3.2, if buyers suffer from imperfect information or salience effects, they may react to transient or foreseen changes in pollution exposure. If this is case, we should see the effect of aermod_{it} vary seasonally, with a peak in Q2 if toxicity is more important and a peak in Q4 if salience is more important. In contrast, if there are no information or salience problems, then we should see aermod_{it} have a similar effect in every quarter, since exposure at any one time should not matter relative to long-term exposure.

Table 1.11 estimates Equation (1.10) which allows the effect of aermod_{it} to vary by quarter. Column 1 reports the 2SLS regression using the annual set of instruments to identify the four endogenous regressors and Column 2 reports the analogous LIML estimates. In both specifications, aermod_{it} has the biggest effect in Q4, consistent with a model where agents use their physical senses to detect pollution and fail to anticipate future pollution exposure when pollution is less salient. However, while the point estimate on Q2 is statistically imprecise, it is still larger than the point estimates in Table 1.5, which may suggest that even though ozone is not easy to see, it is so toxic that the market may still respond do it, if only partially. The fact that the Q2 effect is about half the size of the Q4 effect suggests that visibility effects dominate toxicity. The small or wrong-signed effects in Q1 and Q3, when both salience and toxicity are middling, further support the conclusion that buyers are incorrectly assessing long-run air quality.

Table 1.11: *Exposure's Effect on House Price by Quarter, 2SLS*

	(1)	(2)
Aermod×Q1	-0.005	-0.005
	[0.051]	[0.052]
Aermod×Q2	-0.018	-0.019
	[0.014]	[0.014]
Aermod×Q3	0.015	0.015
	[0.012]	[0.012]
Aermod×Q4	-0.033**	-0.033**
	[0.016]	[0.017]
Method	2SLS	LIML
κ	1	1.0001
Test for Equality		
(p-value)		
Q4=Q1	0.651	0.657
Q4=Q2	0.537	0.541
Q4=Q3	0.063	0.065

Notes: N=41,721. Outcome variable is ln house price. Controls include house effects, year-quarter effects, and local quadratic time trends. Excluded instruments are aermod_pre interacted with year dummies. Standard errors, clustered at 100-m grid, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

1.7 Welfare Implications and Conclusion

This paper provides evidence that clean air has a much higher value than previously believed. The estimated MWTP, \$3,272 per $\mu g/m^3$ of exposure to NO_x emissions, is an order of magnitude larger than past estimates (see Section 1.6.4) and also more in line with the expected health benefits (see Section 1.2). The distinguishing feature of these estimates is that they rely on atmospheric science to determine who is and is not exposed to pollution, while standard estimates do not. When re-estimated using standard, non–wind-based measures, MWTP is small or wrongly signed and statistically insignificant.

Furthermore, the econometric problems behind this difference are not unique to the housing market, raising the concern that other estimates of pollution's effects, like those on infant health, are also biased. This in turn raises the question of whether the MWTP estimated here does indeed cover the estimated health costs since they may be downward biased themselves and is a topic for future work.[37]

The fact that air pollution is far more costly than previously believed has significant policy implications, as air quality regulations are likely to be undervalued. For example, Fowlie, Holland, and Mansur (2012) note that RECLAIM has been frequently criticized as an ineffective policy. But the results here imply that reducing emissions in SCAQMD from 1995 levels to the 2005 RTC cap is worth roughly \$502 million annually, far more than the estimated annual abatement costs of \$38 million.[38] The EPA's troubled attempts to tighten ozone standards, which met resistance on cost-benefit grounds, are another possible example of policy that is grossly undervalued.[39] Optimal subsidies for renewable energy

37. Most estimates of the mortality and morbidity dose response to pollutants are from the epidemiology literature and may suffer from omitted variables bias as well. Thus, it is not immediately clear whether current estimates of direct health effects are too high or too low.

38. There are naturally many general equilibrium costs to consider as well, like those borne by displaced workers (see Walker 2013). SCAQMD asks firms to report how many jobs are lost or gained due to RECLAIM every year. Through 1999, firms reported a total net employment change of −109 workers which they attributed to RECLAIM (SCAQMD 2000). Abatement costs based on SCAQMD calculations (SCAQMD 2000). See also Footnote 2.

39. See, e.g., "Obama Asks EPA to Pull Ozone Rule," *Wall Street Journal*, September 3, 2011; "EPA Sets New Ozone Standard, Disappointing All Sides," *New York Times*, October 1, 2015.

research and electric vehicle take-up are other potential examples.

Vehicle emissions standards are yet another example of potentially undervalued policy. Since coming to light, several back-of-the-envelope estimates of the costs of Volkswagen's cheating on diesel emission tests have been put forward. The Associated Press cite a rough estimate from environmental engineers of $40–170 million per year due to mortality (Borenstein 2015). The radio magazine *Marketplace* cites economists' rough estimates of health costs of $80 million per year (Garrison 2015). Back-of-envelope damages based on this paper's results imply the total cheating cost in the United States alone was $282 million per year.[40]

Consumer welfare is also affected by the fact that people sometimes have a hard time discerning between areas with clean and dirty air, which affects their valuation of homes and where they choose to live. Agents only respond to visible NO_x, not invisible ozone, even though ozone is far more toxic. This could lead to "perverse" sorting, where people with strong preferences for clean air sort into more hazardous areas because ozone-rich air still looks clean. The problem of imperfect information and/or salience could potentially be solved through a cheap informational intervention; providing neighborhood-level information about seasonal and long-term pollution trends for houses on the market could yield large welfare gains per dollar spent.

However, the evidence of sorting found by Chapter 2 suggests that the large aggregate welfare gains disproportionately went to high-income households. This raises the concern that there is steep trade off between equity and efficiency, however large the efficiency gains may be.

40. To get this number, I assume that the extra Volkswagen NO_x emissions were emitted uniformly by SCAQMD firms, then multiply MWTP ($3,272) with the resulting exposure and the households exposed.

Chapter 2

Locational Sorting and the Incidence of Local Amenity Improvements

2.1 Introduction

The equity impacts of air quality regulations are the source of much interest, in part because poor neighborhoods are widely believed to be disproportionately more polluted. The sorting behavior of households is of particular interest because a policy to improve air quality is tied to a location rather than individuals, and sorting makes such policies difficult to implement efficiently (Glaeser and Gottlieb 2008; Kline 2010). The overall progressivity of these policies depends on two related questions: Do rents capitalize air quality as much as housing prices and do poor households get priced out after an air quality improvement?[1]

This paper seeks to answer both of these questions, for which past empirical evidence has been mixed. Bento, Freedman, and Lang (2015) find no change in demographics or population in areas affected by the Clean Air Act Amendments (CAAA) of 1990, and Currie et al. (2015) find no change in the demographic characteristics of mothers in an area where a large polluting firm has recently opened or closed. In contrast, Banzhaf and Walsh (2008)

1. Air quality regulations are thought to be potentially regressive in a number of ways, such as by raising the cost of energy-intensive goods and by changing the capital-labor ratio in the production of these goods (Bento 2013).

and Davis (2011) find that demographics, population, and average income do significantly change with changes in pollution levels. Davis (2011) also finds rents and owner-occupied housing prices respond similarly to the opening of a nearby factory, while Grainger (2012) and Bento, Freedman, and Lang (2015) find rents respond significantly less than prices.

First, I extend the spatial equilibrium model of Epple, Filimon, and Romer (1984) to show that lower-income households will flee air quality improvements if the marginal willingness to pay (MWTP) for clean air increases quickly with income. Epple, Filimon, and Romer's model is a rigorous characterisation of Tiebout's (1956) hypothesis that households sort themselves into communities stratified by income, housing prices, and community-specific amenities. Under this model, communities are defined by their boundaries, which are the incomes of their richest and poorest residents. And while it has been used many times in the context of air pollution (Banzhaf and Walsh 2008) and other amenities like school quality and crime (Epple and Sieg 1999), the focus when deriving predictions has generally been on a world with two communities and a single endogenous boundary between them.

I extend the model to a world with three communities, highlighting the dynamics of the middle community which, with endogenous upper and lower boundaries, more realistically represents the average neighborhood in empirical applications. The model predicts that if this neighborhood's richest residents value the amenity significantly more than its poorest residents, then many of the poorer residents will leave if the amenity is improved.[2] This has the counter-intuitive implication that the poorer residents have an incentive to prevent their own neighborhood from improving and contributes to the more general literature on sorting, amenities, and gentrification.[3] In the specific context of air pollution, it implies that policies that improve local air quality may very well be regressive.

I then estimate the changes in neighborhood rents and demographics due to the air

2. The single-crossing property in prices and the amenity, which is widely used in this literature to guarantee the existence of a stratified equilibrium, already implies that marginal willingness to pay increases with income. The critical factor for whether low-income households will flee amenity improvements is how quickly MWTP increases with income.

3. Avery and Pathak (2015) provide a similar result in the context of school quality with peer effects.

quality improvement in metropolitan Los Angeles caused by the California Electricity Crisis of 2000 and the RECLAIM cap-and-trade program. The Crisis caused an unexpected shortage in RECLAIM's permits for NO_x emissions, inducing firms to install abatement technology, ultimately causing a permanent decrease in exposure to industrial NO_x emissions. I measure local exposure to industrial pollution across metropolitan Los Angeles by using administrative firm-level data on NO_x emissions and an atmospheric dispersion model, AERMOD, which maps firm-level emissions to pollution levels in nearby neighborhoods. I favor this measure over those used in the previous literature because standard measures frequently introduce significant bias, even in quasi-experimental research designs (see Chapter 1).

The results show that rents increased significantly after the air quality improvement with an implied MWTP very close to the sales-based MWTP estimated in Chapter 1. The results also show that improved neighborhoods became richer and better educated but less populous. This decrease in population was driven by the emigration of low-education adults from improved areas. If these leavers were not home owners, and so did not enjoy a wealth windfall from the increase in property values, they were likely made worse off by the amenity improvement. I explore this further by comparing home-ownership rates and price windfalls across the income distribution. The strong correlation between income and home-ownership implies that if the majority of the welfare benefits were captured by changes in property values, then the air quality improvement was strongly regressive.

2.2 Theory

In this section I present a model that closely follows the framework of Epple, Filimon, and Romer (1984) and subsequently Epple and Sieg (1999), Banzhaf and Walsh (2008), and others. Households, defined by their income, sort themselves into communities which are defined by their (exogenous) level of amenities and their (endogenous) housing price. In equilibrium, communities' prices, amenities, and incomes are stratified, with the richest households paying the most for housing and enjoying the highest level of amenities. The

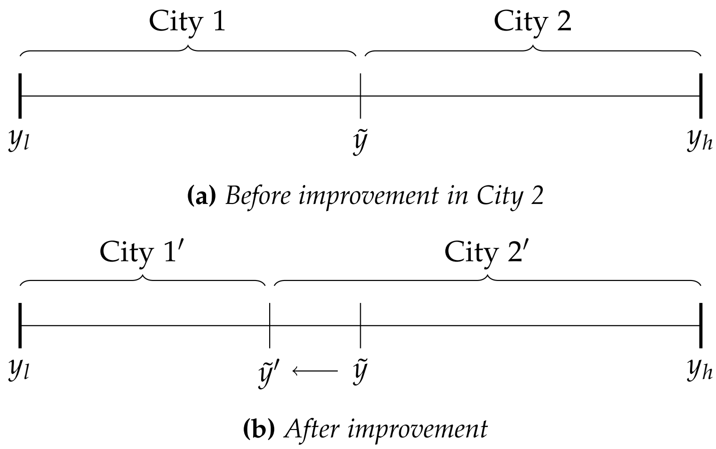

(a) *Before improvement in City 2*

(b) *After improvement*

Figure 2.1: *Shock to a 2-city Equilibrium*

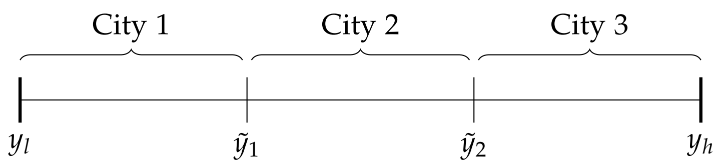

Figure 2.2: *A 3-city Equilibrium*

effects of exogenous changes to amenities on prices and the allocation of households across communities can then be derived from this equilibrium.

The model I present is unchanged from this long line of previous work with a single exception: Where prior work has used a scenario with two communities to derive comparative statics, I use a scenario with three. With two communities stratified by income and amenities, changes to either amenity bundle results in population flowing into the newly improved area exactly as predicted by Tiebout (1956) (see, e.g., Banzhaf and Walsh 2008). This transition is depicted in Figure 2.1. If the higher-income community improves its amenities, the previously marginal households with income $\tilde{y}$ will move into the newly improved community.

However, the dynamics of a third community in the middle of the income distribution are more complex because both of its boundaries are endogenous (see Figure 2.2). If the middle community, City 2, improves its amenities, the initial impetus will be for marginal households at both ends ($\tilde{y}_1$ and $\tilde{y}_2$) to flow into City 2. But if higher-income households have a higher marginal willingness to pay (MWTP) for the amenity, they will bid up prices more than lower-income households would have on their own, stifling the flow of lower-income immigrants. If the difference in MWTP is large enough, there may be no net migration from the bottom of the income distribution or even negative net migration, with low-income incumbents being pushed out. This scenario is not possible in a two-city model.

Including a third city highlights the asymmetric behavior of households created by a MWTP function that increases with income. This has significant implications for the incidence of place-based policies. Most significantly, it suggests that, all else equal, such policies will tend to be regressive in nature. In addition, most neighborhoods in empirical applications will be "middle" communities since for most neighborhoods there are both better and worse alternatives.

2.2.1 Model Setup

Definition 1 (Community)**.** A community $j \in \{1, \ldots, J\}$ is characterized by (p_j, g_j), its endogenous unit price of housing and its exogenous amenity level. Housing in community j is supplied according to $S^j(p)$ which has the following properties: $S^j_p > 0$; $S^j(p^l_j) = 0$ for some lower bounding price $p^l_j > 0$; and $0 < S^j(p) < \infty$ for $p > p^l_j$.

Definition 2 (Household)**.** A household is characterized by its income, y, which follows a distribution $f(y)$ with continuous support $[y_l, y_h]$. A household's preferences are characterized by indirect utility $V(y, p, g)$ which is assumed to have the following basic properties: $V_y > 0$, $V_p < 0$, $V_g > 0$. Households also have housing demand $h(p, y)$ which is assumed to be independent of g and to have the following properties: $h_p < 0$, $h_y > 0$, $0 < h < \infty$.[4]

Households maximize their utility by choosing the community (p_j, g_j) that maximizes their utility. An equilibrium will be a set of prices and household-community pairings where no household wants to move to a different community. Generally speaking, characterizing

4. See Epple and Sieg (1999) or Sieg et al. (2004) for an example and discussion of an indirect utility function that satisfies these properties.

this equilibrium requires some additional structure. A common approach, and the one I follow here, is to assume V satisfies single-crossing in p and g.[5]

Assumption 1 (Single-crossing preferences)**.** Assume $V(y,p,g)$ satisfies the *single-crossing property*. That is, assume

$$M(y,p,g) = -\frac{V_g(y,p,g)}{V_p(y,p,g)} = \left.\frac{\mathrm{d}p}{\mathrm{d}g}\right|_{V=\bar{V}(y)}$$

is monotonically increasing in y for all (y, p, g).

For a fixed y, $M(y,p,g)$ is the slope of an indifference curve in (p,g) space. Assumption 1 requires that the slope of this curve be increasing with y. It is important to note that $M(y,p,g)$ is also the MWTP for g at point (y,p,g); it represents how much more the household is willing to pay (dp) for a marginal change in the amenity (dg). Thus, another way to frame Assumption 1 is that MWTP must be increasing in y. It is this increasing MWTP with income that will drive the asymmetric behavior of high- and low-income households.

Epple, Filimon, and Romer (1984) show that with Assumption 1, an equilibrium where no household wants to move must be of the following type.[6]

Definition 3 (Equilibrium)**.** For ordered communities $g_j < g_{j+1}$, an equilibrium is a set of incomes $\{\tilde{y}_1, \ldots, \tilde{y}_{J-1}\}$ and prices $\{p_1, \ldots, p_J\}$ such that $\tilde{y}_j < \tilde{y}_{j+1}$ and $p_j < p_{j+1}$, where the following conditions hold:

$$V\left(\tilde{y}_j, p_j, g_j\right) = V\left(\tilde{y}_j, p_{j+1}, g_{j+1}\right) \quad \forall j < J \tag{2.1}$$

$$\int_{\tilde{y}_{j-1}}^{\tilde{y}_j} h(p_j, y) f(y)\,\mathrm{d}y = S^j(p_j) \quad \forall j \tag{2.2}$$

where $\tilde{y}_0 = y_l$ and $\tilde{y}_J = y_h$.

5. This is the approach taken by, e.g., Epple and Sieg (1999), Sieg et al. (2004), Banzhaf and Walsh (2008), and Banzhaf (2015). See Kuminoff, Smith, and Timmins (2013) for a review.

6. Epple and Sieg (1999) extend this model to include heterogeneity in preferences as well as income. They show that the equilibrium is similarly characterized, with community boundaries $\tilde{y}_j(\alpha)$ being lines in income-preference space.

The first set of conditions are boundary indifference conditions, which require the household which borders j and $j+1$ to be indifferent between the two. The second set of conditions are internal equilibrium conditions for each community, which require housing demanded to equal housing supplied. Epple, Filimon, and Romer (1993) show that such an equilibrium exists and is unique.

2.2.2 Predictions

When deriving comparative statics, past work has generally used a model with two communities ($J = 2$).[7] This is likely because there are $2J-1$ equilibrium conditions—$J-1$ cross-community spatial equilibrium conditions and J within-community supply-demand conditions—and derivations quickly become intractable as J grows. Unfortunately, the imposition of $J = 2$ implicitly assumes that all households only have one alternative to their current neighborhood: low-city residents can move up and high-city residents can move down. More technically speaking, this imposes all communities have only one endogenous boundary.

However, when $J \geq 3$, at least one community has two endogenous boundaries and residents can move up *or* down the amenity ladder. This creates an opportunity for incumbent types to flee amenity improvements. The propositions that follow provide conditions under which this may occur. Proofs of all propositions are given in Section B.1.

Suppose $J = 3$, so the vector of endogenous variables is $(\tilde{y}_1, \tilde{y}_2, p_1, p_2, p_3)$ and the exogenous variables of interest are (g_1, g_2, g_3).

Proposition 1. *The following conditions hold:*

$$\frac{\partial \tilde{y}_2}{\partial g_2} > 0; \quad \frac{\partial p_2}{\partial g_2} > 0; \quad \frac{\partial p_3}{\partial g_2} < 0 \tag{2.3}$$

$$\frac{\partial \tilde{y}_1}{\partial g_2} \propto \frac{\partial p_1}{\partial g_2} \tag{2.4}$$

7. For example, Banzhaf and Walsh (2008) show that the comparative statics in the two-city case align exactly with the intuition of Tiebout (1956), with population flowing into areas that experience an increase in amenities (see Figure 2.1).

and the sign of $\partial \tilde{y}_1 / \partial g_2$ is ambiguous.

The signed comparative statics are not surprising. In response to the improvement to Community 2, the price of housing there increases and some residents of Community 3 move in. The proportionality in Equation (2.4) also follows economic intuition, with p_1 decreasing if some Community 1 residents move up to Community 2 and p_1 increasing if some residents of Community 2 flee into Community 1. The fact that these last two effects cannot be signed without additional conditions is consistent with the previously discussed intuition.

Proposition 2 gives the necessary and sufficient condition for low-income residents in Community 2 to flee in response to improvements in g. Once again consistent with intuition, this condition hinges on the difference in MWTP between high- and low-income households.

Proposition 2.

$$\frac{\partial \tilde{y}_1}{\partial g_2} > 0$$

if and only if $M(\tilde{y}_2, p_2, g_2)$ is sufficiently large relative to $M(\tilde{y}_1, p_2, g_2)$ such that

$$\frac{M(\tilde{y}_2, p_2, g_2)}{M(\tilde{y}_1, p_2, g_2)} - 1 > \frac{H^3 \dfrac{\partial V^3}{\partial \tilde{H}^3} + V_y(\tilde{y}_2, p_3, g_3) - V_y(\tilde{y}_2, p_2, g_2)}{H^2 \dfrac{\partial V^2}{\partial \tilde{H}^2}} \tag{2.5}$$

where

$$\frac{\partial V^j}{\partial \tilde{H}^j} = \frac{V_p(\tilde{y}_2, p_j, g_j)}{\int_{\tilde{y}_j}^{\tilde{y}_{j+1}} h_p(p_j, y) f(y)\, \mathrm{d}y - S_p^j(p_j)} \tag{2.6}$$

and $H^j = h(p_j, \tilde{y}_2) f(\tilde{y}_2)$

The left side of the inequality is the percentage difference in MWTP of $\tilde{y}_2$ and $\tilde{y}_1$, the richest and poorest residents of Community 2. The right side of the inequality captures the trade-offs households at $\tilde{y}_2$ face when choosing between Community 2 and Community 3.

First, consider Equation (2.6), which is the change in $\tilde{y}_2$'s utility in city j per unit of excess housing demand. The denominator is the derivative with respect to price of excess housing demand in j, i.e., the gap between the supply and demand curves for prices just above the equilibrium price. It can also be thought of as the rate at which vacant housing units are created by an infinitesimal change in price. The numerator is simply the change in $\tilde{y}_2$'s utility in j due to a change in p_j. If we write excess demand as

$$\tilde{H}_j = \int_{\tilde{y}_j}^{\tilde{y}_{j+1}} h(p_j, y) f(y)\, \mathrm{d}y - S^j(p_j)$$

and abuse some notation, we see that

$$\frac{\frac{\partial V(\tilde{y}_2, p_j, g_j)}{\partial p}}{\frac{\partial \tilde{H}_j}{\partial p}} = \frac{\partial V^j}{\partial p} \cdot \frac{\partial p}{\partial \tilde{H}_j} = \frac{\partial V^j}{\partial \tilde{H}^j}$$

This can also be thought of as the rate at which increased demand for housing in j changes $V(\tilde{y}_2, p_j, g_j)$. When multiplied by H^j, it is the amount $\tilde{y}_2$ households change their utility by increasing housing demand in j.

The V_y terms in Equation (2.6) capture the substitution effect faced by $\tilde{y}_2$ between p and g in the two communities. This can be seen more clearly by noting

$$\begin{aligned} &V_y(\tilde{y}_2, p_3, g_3) - V_y(\tilde{y}_2, p_2, g_2) \\ &= V_y(\tilde{y}_2, p_3, g_3) - V_y(\tilde{y}_2, p_3, g_2) + V_y(\tilde{y}_2, p_3, g_2) - V_y(\tilde{y}_2, p_2, g_2) \\ &\approx (g_3 - g_2)\, \frac{\partial V_y(\tilde{y}_2, p_3, g_3)}{\partial g} + (p_3 - p_2)\, \frac{\partial V_y(\tilde{y}_2, p_2, g_2)}{\partial p} \end{aligned}$$

Together, the right side of Equation (2.5) is the housing- and amenity-related utility effects of choosing Community 3, normalized by the housing-related utility effect of choosing Community 2.

Finally, we see that the comparative statics with respect to g_1 and g_3 are as expected.

Proposition 3. *The comparative statics for g_1 are g_3 are*

$$\frac{\partial \tilde{y}_1}{\partial g_1} > 0; \quad \frac{\partial \tilde{y}_2}{\partial g_1} > 0; \quad \frac{\partial p_1}{\partial g_1} > 0; \quad \frac{\partial p_2}{\partial g_1} < 0; \quad \frac{\partial p_3}{\partial g_1} < 0;$$

$$\frac{\partial \tilde{y}_1}{\partial g_3} < 0; \quad \frac{\partial \tilde{y}_2}{\partial g_3} < 0; \quad \frac{\partial p_1}{\partial g_3} < 0; \quad \frac{\partial p_2}{\partial g_3} < 0; \quad \frac{\partial p_3}{\partial g_3} > 0;$$

2.2.3 Discussion

The possibility of emigrants from a middle city is important for welfare because it makes *any* amenity improvement more regressive in a number of ways.

First and foremost, displaced renters are made unambiguously worse off by improvements in their own community.[8] At the same time, many home-owners will benefit from the amenity improvement even if they decide to leave because of they also enjoyed a positive wealth shock through the home's appreciation. This makes the correlation between income and home-ownership important for understanding incidence, and I explore this correlation empirically in Section 2.5.4.

Second, the existence of emigrants creates a negative pecuniary externality on the emigrants' new community. Since $p_1' > p_1$, incumbent households in Community 1 are worse off due to improvements in Community 2. However, this welfare effect is likely to be small relative to the effect on emigrants themselves as long as the emigrants disperse to many different destination communities.

Before testing the model empirically, it should be noted that two of its simplifying assumptions are certain to conflict with the data in predictable ways. The first assumption is heterogeneity of preferences. Epple and Sieg (1999) extend the core model presented in Section 2.2.1 to include heterogeneity of preferences as well as income. For an amenity preference parameter α, the set of boundary households between cities j and $j+1$ is $\tilde{y}_j(\alpha)$, a line in (α, y) space. The second assumption is that households value only one spatial amenity, g_j. In actuality, g_j will be a bundle of amenities, which will not always themselves be stratified. For example, a central business district may be polluted and noisy but offer short commute times and easy access to restaurants, museums, etc.

Relaxing these assumptions allows for high-income households that choose high-pollution neighborhoods, which we will see evidence of in Section 2.5.3. However, conditional on preferences and other spatial amenities, the predictions made above remain the

8. This follows from a simple revealed preference argument. Suppose g_2 exogenously changes to $g_2' > g_2$ and the condition for Proposition 2 holds, so $\tilde{y}_1' > \tilde{y}_1$ and $p_1' > p_1$. By definition of $\tilde{y}_1$ and $V_p < 0$, we know $V(y, p_1', g_1) < V(y, p_1, g_1) < V(y, p_2, g_2)$ for $y > \tilde{y}_1$.

same.

2.3 Research Design

This section describes the key features of the research design. First, how air quality g_j is measured for each community using the atmospheric dispersion model AERMOD. Next, how the California Electricity Crisis created an exogenous shock to the air quality of neighborhoods in greater Los Angeles. And finally, how I use the Crisis to measure household sorting behavior after the amenity shock.

2.3.1 Measuring Air Quality with AERMOD

In order to measure the effect of air quality on neighborhood characteristics, you need data on neighborhood-level air quality. The predominant method for measuring local air quality is to use pollution monitors and interpolate to the centroid of the geographic unit of interest.[9] Distance from a pollution source is another possible metric.

However, these methods are too coarse to capture the sudden changes in pollution levels across space, especially around pollution sources themselves. As Chapter 1 shows, this biases estimates toward zero, often severely. Thus, I follow the methodology of Chapter 1 and use AERMOD, an atmospheric dispersion model, to measure local exposure. AERMOD uses detailed data on local meteorology and pollution sources to map the impact of individual firms' emissions on any arbitrary nearby location. Using AERMOD, I construct aermod_{nt}, measure of block group n's exposure to industrial NO_x emissions at time t.[10]

2.3.2 Electricity Crisis as Natural Experiment

The natural experiment used is the same as in Section 1.4.1. The variable $\text{aermod\_pre}_n \times \text{post}_t$, where $\text{post}_t = \mathbf{1}\{y \geq 2001\}$, can then be used an instrument which captures the

9. In the pollution-health literature, using ZIP codes and their centroids is especially common. See Currie and Neidell (2005), Schlenker and Walker (Forthcoming), and Knittel, Miller, and Sanders (2014), among others.

10. See Section 2.4.2 for a complete discussion of how this is done.

differential effect of the Crisis on neighborhood n. This instrument is the equivalent of a difference-in-difference estimate with variable treatment intensity.

The identification assumptions remain the same (again see Section 1.4.1). One way to assess the identification assumption is to plot the effect of the instrument over time on both the first stage (pollution exposure) and the outcome of interest. If these event studies show no trend before the exogenous shock, followed by a sharp change in outcomes immediately after the shock, this is strong evidence that the common trends assumption holds. Unfortunately, the block group data available for this paper includes only two observations across time for each block group, making such an event study impossible. However, Chapter 1, which looks at the effect of the Crisis on quarterly house prices, provides an event study, replicated in Figure 1.6. This figure shows a sharp break in both pollution exposure and house prices that coincides with the Crisis, suggesting that the Crisis makes a good natural experiment.

2.3.3 Estimation Strategy

The following empirical equation relates block group SES outcomes to pollution exposure:

$$y_{nt} = \text{aermod}_{nt}\beta + \text{post}_t\alpha + \delta_n + \left(\boldsymbol{W}_n \times \text{post}_t\right)\boldsymbol{\Gamma} + \varepsilon_{nt} \tag{2.7}$$

where y_{nt} is some characteristic of block group n (e.g., population) in period t; aermod_{nt} is exposure to industrial NO_x-based pollution; δ_n are block group fixed effects; $\boldsymbol{W}_n$ is a vector of time-invariant block group characteristics detailed below; and ε_{nt} is the usual residual term. Period t indexes data from either the 2000 Census, before the Crisis, or the ACS 5-year average for 2005–2009, the earliest available block group data from after the Crisis. These controls account for a number of factors that may confound estimates of β. The block group fixed effects, δ_n, capture of all time-invariant characteristics about the neighborhood.

The vector $\boldsymbol{W}_n$ controls for differential effects by block group characteristics over time by including two kinds variables. The first kind is dummy variables for the block group's location within the metropolitan area, defined by a 10 km grid.[11] This allows different parts of the metropolitan area to have different secular trends. The second kind of variables in $\boldsymbol{W}_n$ control for the block group's socio-economic characteristics in the year 2000. For example, if labor market shocks over this period differentially affect low- and high-education workers, this could be captured in part by β if pollution exposure is correlated with educational attainment. To solve this problem, I include the following year 2000 block group characteristics in $\boldsymbol{W}_n$: population; number of households; log median household income; population over age 25; fraction of adults over age 25 with no high school diploma; fraction with high school diploma but no time at college; fraction white (non-Hispanic); fraction Hispanic; fraction black. These variables are discussed further in Section 2.4.1.

Using the instrument described in Section 2.3.2, the reduced-form estimate of the Crisis' effect is simply

$$y_{nt} = \left(\text{aermod\_pre}_n \times \text{post}_t\right)\pi + \text{post}_t\alpha + \delta_n + \left(\boldsymbol{W}_n \times \text{post}_t\right)\boldsymbol{\Gamma} + \varepsilon_{nt} \tag{2.8}$$

The model presented in Section 2.2 stresses stratification in income, with households across the income distribution responding differently. Estimating differential behavior by income over time is problematic when using aggregated data, primarily because neighborhoods are not homogeneous. While it may be intuitive to use median income to characterize a block group, we are also interested in the behavior of relatively lower- and higher-income households within each block group. I attempt to solve this by characterizing block groups by the educational attainment of their residents who are over age 25. Those without a high school diploma are "low" education or "less than high school"; those with only a high school diploma are "middle" ed; and those who have had at least some college are "high" ed or "more than high school." I can then estimate variations of Equation (2.8) that allow π

11. Given the large size of the sample region, the ideal geographic unit for these trends would be individual cities, which have economically meaningful boundaries (unlike zip codes) and are generally small but not so small as to be computationally burdensome (unlike tracts and zip codes). Unfortunately, many cities are not geographically convex, and the cities of Los Angeles and Long Beach cover a large portion of the sample region while also having a great deal of within-city heterogeneity. To overcome these issues, I use a 10-km grid, shown in Figure A.1, which is aligned to preserve as many city boundaries as possible. This grid results in 17 different areas that each get their own time effects.

to vary with the year 2000 educational composition of the block group.

Section 2.2.2 also makes several strong predictions about how the cost of housing responds to amenity changes. While measures of home value and monthly rent appear in the Census and ACS, these data can be problematic (see Section 2.4.1). To supplement the block group–level data on rents, I also estimate a handful of specifications using house-level sales data, borrowing from Chapter 1.

$$\ln p_{it} = \text{aermod}_{it}\beta + \alpha_i + \delta_t + (\boldsymbol{W}_i \times t)\,\boldsymbol{\Gamma}_1 + \left(\boldsymbol{W}_i \times t^2\right)\boldsymbol{\Gamma}_2 + \varepsilon_{it} \tag{2.9}$$

where i is an individual house and t indexes quarters (e.g., 1999Q3). The controls in $\boldsymbol{W}_i$ are similar to those in Equation (2.7). First, dummy variables for local 10 km area as described above. Second, controls for the neighborhood's socio-economic makeup: the average loan-to-value ratio for houses sold in the house's census tract in 2000; the average predicted interest rate for mortgages taken out in the house's census tract in 2000; the median household income in the house's census block group in 2000; fraction of adults over age 25 with no high school diploma; and fraction with high school diploma but no time at college. The time period of the house sample is 1997–2005.

I restrict the region of analysis to the southwest region of SCAQMD territory, roughly between Santa Monica and Huntington Beach (see Figure 1.3a), to minimize measurement error due to geography. All of the major polluters are located in this region and locations farther away from the pollution sources are likely to have less actual exposure from the firms and more noise in the modeling prediction, decreasing the signal-to-noise ratio of the pollution measure. Predicting the pollution distribution is also more complicated farther inland because of the San Gabriel and Santa Ana Mountains, which can act like a dam, collecting pollution blown from the coasts. To avoid these problems, I restrict my sample to houses within 10 kilometers of a major electric firm in Los Angeles or Orange County.[12]

12. I also include in this group the southwestern most firm in the area in order to include the Palos Verdes Peninsula in the regression sample (see Figure A.1).

2.4 Data

Data for houses, firms, and meteorology are the same as in Section 1.5.

2.4.1 Census Block Groups

Data on Census block group demographics are taken from the 2000 Census and 2005–2009 5-year American Community Survey (ACS) sample. For each block group, these data include total population; white (non-Hispanic) population; Hispanic population; black population; the number of households; median of household income; median rent; and educational attainment for individuals age 25 and older. The data also include the block groups' total land area, which I use to calculate population density (population per square mile). I group educational attainment into three categories: people who did not graduate high school; people who graduated high school but do not have a bachelor's degree; and people who hold at least a bachelor's degree. To reduce noise, I drop block groups that have less than 400 people in 2000, which is roughly the 4th percentile of all block groups and constitutes less than 0.5% of all people in the sample. In specifications using median rent, I drop observations with top coded values ($2001) in either year. Table 2.1 presents summary statistics for both 2000 and 2005–2009.

2.4.2 AERMOD-based Measure of Exposure

I use AERMOD, which maps firm-level output to local exposure, to construct a measure of a block group's exposure from all industrial sources. Software for using AERMOD is available on the EPA's website and includes documentation, Fortran source code, and pre-compiled executables for Windows.[13]

Location ℓ's exposure to $\mathrm{NO_x}$ emissions from firm f at time t can be written $\mathrm{NOx}_{ft} \cdot h(d_{f\ell}, \theta_{f\ell}; \boldsymbol{S}_f)$, where $\boldsymbol{S}_f$ contains information on the firm's smoke stacks, as well as local

13. See http://www.epa.gov/scram001/dispersion_prefrec.htm. I use AERMOD version 13350, compiled using Intel Fortran Compiler 15.0 for Linux and run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University.

Table 2.1: *Block Group Summary Statistics*

	Total		Mean	
	2000	2005/9	2000	2005/9
Population	2,775,700	2,811,468	1,435	1,454
			(814)	(867)
Households	950,591	952,008	492	492
			(322)	(332)
Pop. Density (pop/mi^2)			13,423	13,518
			(8,389)	(8,815)
Household Income (BG Median)			49,292	64,211
			(23,411)	(32,920)
Population over age 25	1,717,881	1,796,814	888	929
			(505)	(564)
Educational Attainment (count)				
Less than High School	458,399	384,055	237	199
			(221)	(209)
High School Grad	830,050	895,603	429	463
			(269)	(294)
More than High School	429,432	517,156	222	267
			(262)	(312)
Educational Attainment (fraction)				
Less than High School			0.28	0.22
			(0.22)	(0.19)
High School Grad			0.48	0.50
			(0.13)	(0.14)
More than High School			0.24	0.28
			(0.19)	(0.21)
Race/Ethnicity (count)				
White (non-Hispanic)	852,136	787,815	441	407
			(468)	(466)
Hispanic	1,030,236	1,147,634	533	593
			(546)	(601)
Black	507,488	468,462	262	242
			(380)	(378)
Race/Ethnicity (fraction)				
White (non-Hispanic)			0.34	0.32
			(0.31)	(0.30)
Hispanic			0.34	0.38
			(0.26)	(0.28)
Black			0.19	0.17
			(0.25)	(0.24)

Notes: Number of block groups is 1,934. Block groups with fewer than 400 people in 2000 are excluded from regression sample and so are excluded here. Data for 2000 comes from the 2000 Census. Data for 2005/9 comes from the 2005–2009 ACS 5-year sample and is labeled “2005” elsewhere. All educational attainment variables are restricted to people who are at least 25 years old. Income is denominated in nominal dollars. Standard deviations in parentheses.

meteorological conditions. The data I use for NOx_{ft} and $\boldsymbol{S}_f$ are described in Sections 1.5.2 and 1.5.3. A firm's meteorological data is taken from the nearest meteorology monitor. The values for $(d_{f\ell}, \theta_{f\ell})$ are calculated by AERMOD from firms and houses' latitude and longitude. AERMOD then outputs $\text{aermod}_{\ell f t}$, the location's exposure to the firm's emissions. The location's total exposure to industrial NO_x emissions is simply $\text{aermod}_{\ell t} = \sum_f \text{aermod}_{\ell f t}$.

For block group–level exposure, I first calculate exposure at the block level, then calculate the population-weighted average for each block group. At the block level, I use the process described above where ℓ is the Census-provided internal point for each block.[14] This is a more attractive approach than using the block group's internal point because it accounts for heterogeneity in population and exposure across the block group and is a closer approximation to the average exposure to the block group's residents. For house-level exposure, I use

Because AERMOD loops over all firms, locations, and meteorological data, it is very computationally intensive for such a large sample, so I impose several restrictions on the data to make calculation more feasible.[15] First, I only calculate a firm's exposure to houses that are within 20 kilometers of the firm and set exposure outside this radius to zero. Second, I use one year of meteorological data, 2009, which is also the only year during which all of the meteorological stations described in Section 1.5.3 were operating. Third, for houses, I round each house's latitude and longitude coordinates to the nearest 100 meters and assign houses within the same 100-meter grid square the same exposure.

14. Analyses using Census geographies like block groups or ZCTA's often use the "centroid" of the geography as its the representative point in space. However, the Census Bureau is very particular to note that because these geographies are not convex, the true centroid may lie outside the geography of interest. As a solution, the Census Bureau calculates "internal points," which are constrained to be inside the geography.

15. Even with these restrictions, the model takes approximately 210 CPU days to process all the data.

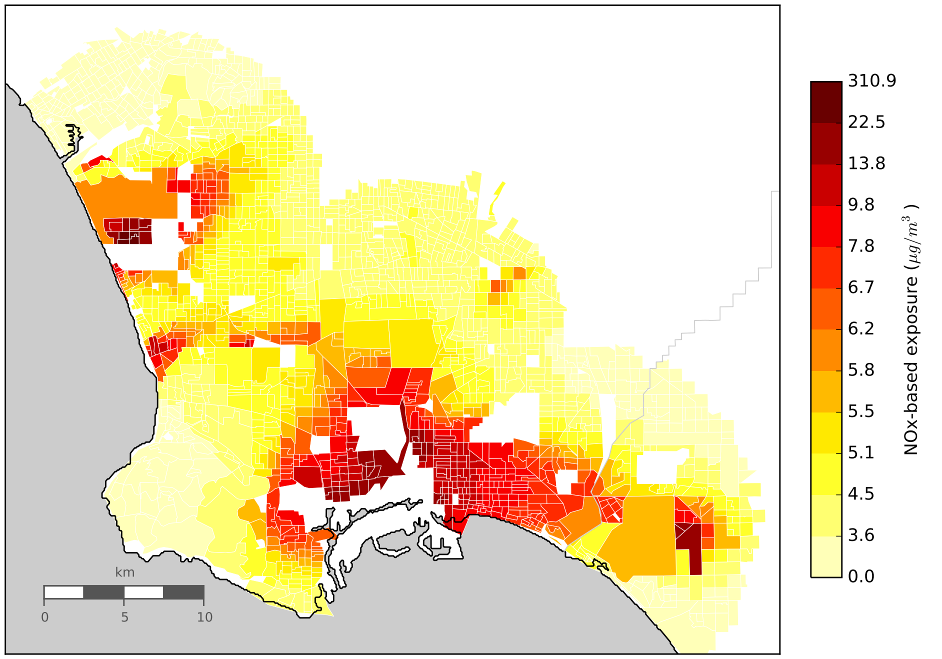

Figure 2.3: *Total Exposure in Sample Region by Block Group, 2000*

Notes: Breaks in the color scale are set at order statistics of the plotted sample in Figure 1.3.

2.5 Empirical Results

2.5.1 Income Stratification

The model predicts stratification by income and amenities. Figure 1.3 shows exposure to industrial NO_x emissions across the sample area while Figure 2.4 shows the fraction of each block group's adults that did not graduate high school. I will refer to these individuals as the "low" education group. Note that in both block group maps, block groups with less than 400 residents are not plotted.

There are many areas where high pollution concentration and low educational attainment are correlated, most notably in the south half of the map, from Palos Verdes in the west, through Long Beach in the center, to Westminster in the east (see Figure A.1 for city and neighborhood names). It is also notable that many of the most polluted areas according to Figure 2.4 are non-residential. This is especially true near Long Beach and Los Angeles

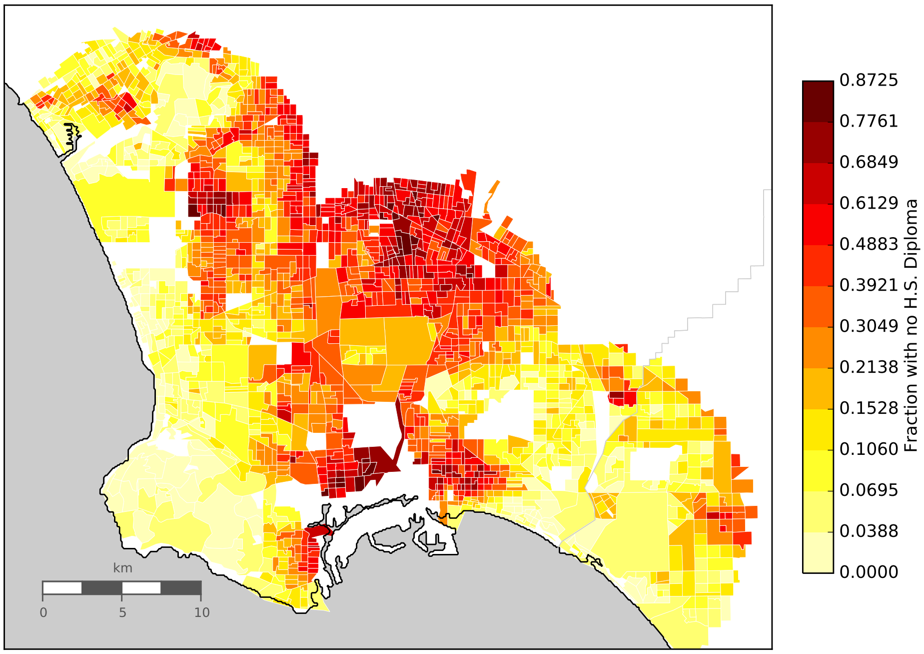

Figure 2.4: *Fraction of Block Group without a High School Diploma*

International Airport.

However, there are also areas where the correlation does not hold. For example, El Segundo and Hermosa Beach in the west are heavily polluted but have high educational attainment, while in Watts, Lynwood, and Compton the opposite is true. This apparent contradiction is likely due to other amenities offered in the area, as discussed in Section 2.2.3. The former cities have direct access to high quality beaches, while the latter cities are well-known for high crime rates and gang violence.

In addition to the correlation of amenities and income, stratification also suggests that neighborhoods would be at least somewhat segregated by income (preference heterogeneity will counteract the raw income segregation to some degree). While Figure 2.4 shows that low-education neighborhoods tend to cluster, it says nothing about the other residents of these neighborhoods. To show that neighborhoods are generally segregated by education level, I define two more education groups: those with only a high school diploma ("middle") and those who have had at some college ("high"). Figure 2.5 is a simplex plot of each block group's fraction of low-, middle-, and high-education residents. For example, a point in

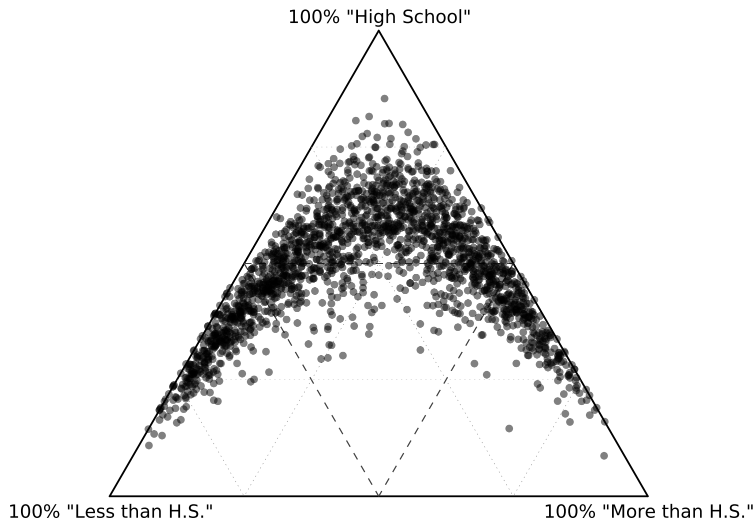

Figure 2.5: *Simplex Scatterplot of Block Group Education*

Notes: Each dot represents a block group's educational attainment. For example, a dot at the top vertex represents a block group where all adults have a high school diploma but did not go to college, while a dot in the middle represents a block group with an equal number of people in each education group. Definitions of education groups are given in Section 2.4.1.

the left corner of the simplex represents a block group that is 100% low education, while a block group in the exact center of the plot has one-third of its population in each group. The dashed lines mark 50% concentrations, e.g., points above the top dashed line are at least 50% middle-ed and points within the dashed triangle have no majority group.

Figure 2.5 shows that households are fairly segregated by educational attainment. Under complete segregation, where any neighborhood has only one education level, the plot would have masses at each vertex. Complete integration, where educational attainment is random across space, would have a mass in the center of the plot. A plot of complete stratification but not segregation would have all points on the left and right edges of the simplex; i.e., low-ed and middle-ed households would sometimes mix, and middle- and high-ed households would sometimes mix, but low- and high-ed households would never mix. The results in Figure 2.5 are most similar to this scenario, as there is very little mixing of low- and high-ed households.

Table 2.2: *Effect on Block Group Median Monthly Rent*

	(1)	(2)	(3)	(4)	(5)	(6)
Aermod_pre×post	0.0025	0.0031*	0.0031**			
	[0.0015]	[0.0018]	[0.0015]			
Aermod				-0.0099	-0.0124	-0.0126**
				[0.0062]	[0.0076]	[0.0063]
Method	OLS	OLS	OLS	2SLS	2SLS	2SLS
Weighted by.		Pop.	Renters		Pop.	Renters
R^2	0.9160	0.9331	0.9536			

Notes: N=3,162. Outcome is log of median rent. Excluded instrument in 2SLS regressions is aermod_pre × post. Rents with error codes ($0) or top codes ($2,001) are dropped from the sample. Sample and controls are otherwise the same as in Table 2.3, plus an interaction of median rent in 2000 with post. Sample average of aermod_pre is 6.781. Standard errors, clustered by tract, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.5.2 Rent

The model in Section 2.2 predicts that housing costs will increase after an amenity improvement. Using the same research design as this chapter, Chapter 1 estimates the effect of the Crisis on house prices using repeat sales. That paper finds that the reduced form effect of the instrument on prices is 0.0032, meaning that each unit of treatment intensity increased home prices by roughly 0.3 percent. The second-stage semi-elasticity of exposure to prices is -0.0073, implying a MWTP to reduce pollution of $3,272.

Table 2.2 shows that rents respond to air pollution in very similar manner. Columns 1–3 shows the reduced-form estimates, based on Equation (2.8), with the natural log of median rent as the dependent variable. The preferred specification is estimated in column 3 which weights each block group by its number of renter households in 2000 so that outliers from block groups with few renters do not skew the results. However, as column 1 (no weights) and column 2 (raw population weights) show, the weights used do not dramatically change the point estimates. The estimates show that log rents increased by 0.0031 for every unit of treatment intensity, very similar to the effect on log house prices (0.0032) from Chapter 1. The second-stage estimates in columns 4–6 are likewise very similar.

This stands in contrast with some of the more recent work on capitalization by rents

Table 2.3: *Effect of Pollution on Block Group Demographics*

	ln Income	% No HS	ln Pop.	ln H-holds	ln H. Units
			A. Naïve OLS		
Aermod	-0.0940***	0.0235***	0.0025	0.0048***	-0.0003
	[0.0079]	[0.0023]	[0.0018]	[0.0014]	[0.0015]
			B. OLS with Controls		
Aermod	0.0019	0.0012	-0.0010	0.0014	0.0060
	[0.0050]	[0.0013]	[0.0053]	[0.0043]	[0.0040]
			C. Reduced Form		
Aermod_pre×post	0.0036	-0.0020***	-0.0027**	-0.0024*	-0.0031**
	[0.0026]	[0.0006]	[0.0014]	[0.0014]	[0.0015]

Notes: N=3,868. Sample periods are 2000 and 2005–2009 using data from the 2000 Census and 2005–2009 ACS, respectively. Regressions include block group fixed effects and 10-km grid–post dummies. Year-2000 demographic controls, interacted with "post", include: population, number of households, number of housing units, ln median household income, number of people at least 25 years old, fraction without a high school diploma, fraction with diploma but no college, fraction white (non-Hispanic), fraction Hispanic, fraction black. All educational attainment variables are restricted to the sample of people who are at least 25 years old. Block groups with fewer than 400 people in 2000 are dropped. Sample average of aermod_pre is 6.560. Standard errors, clustered by tract, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

versus house prices Grainger (2012) and Bento, Freedman, and Lang (2015). It also suggests that low-income households who are predominantly renters are not shielded from the price effects of air quality improvement.

2.5.3 Demographics

Table 2.3 explores how the composition of neighborhoods changes, specifically looking at log of median household income; fraction of adults with no high school diploma; log population; log number of households; and log housing units. The naïve univariate estimates of each outcome on pollution exposure, shown in Sub-table A, are consistent with intuition on sorting. More polluted neighborhoods are poorer and less educated. However, with Equation (2.7)'s full set of controls, estimated in Sub-table B, these correlations become small and statistically insignificant. This suggests that other local conditions and amenities are also important and confirms that an exogenous shock to pollution is needed to identify sorting behavior.

Table 2.4: *Change in Population by Educational Attainment*

	(1)	(2)	(3)	(4)	(5)	(6)
	log Less than HS		log High School		log More than HS	
Aermod_pre×post	-0.021***		0.005**		-0.002	
	[0.005]		[0.002]		[0.004]	
Aermod_pre×post×						
% Less than HS in 2000		-0.0158		0.0196**		-0.0252
		[0.0250]		[0.0084]		[0.0222]
% High School in 2000		-0.0243		-0.0149		0.0747***
		[0.0284]		[0.0103]		[0.0221]
% More than HS in 2000		-0.0204		0.0234*		-0.1025***
		[0.0343]		[0.0128]		[0.0266]
R^2	0.936	0.936	0.938	0.938	0.937	0.938
N	3,592	3,592	3,868	3,868	3,718	3,718

Notes: Outcome is the log of the number of people with the given educational attainment who are at least 25 years old. Observations weighted by total population in 2000. Block groups with an undefined logarithm in either year are dropped. Otherwise, sample and controls are the same as in Table 2.3. Sample average of aermod_pre is 6.560. Standard errors, clustered by tract, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Sub-table C, which estimates the reduced-form effect of the Crisis, reveals a significant sorting response; however, this response does not perfectly mesh with classic Tiebout sorting. As expected, the average income and education attainment in improved neighborhoods increases, though the effect on income is not statistically significant. However, the number of people and households in improved areas *decreases* relative to the counterfactual, with an accompanying decrease in the number of housing units. This is consistent with households fleeing the amenity improvement, though it is not clear who is fleeing.

Table 2.4 estimates the net change in population by educational attainment and shows that low-education households are driving the decrease in population, consistent with the model's predictions. The table separates the population effect by low-education (columns 1 and 2), middle-education (3 and 4) and high-education (5 and 6). Column 1 shows that block groups had 2.1% fewer low-education residents for every unit of treatment intensity. At the average treatment intensity, this implies the Crisis caused a relative decrease in the adult low-ed population of 14%, or 63,000 people. Column 3 shows a small gain of middle-education individuals while Column 5 shows no significant change in high-ed adults. However, because the unit of observation is the block group, these regressions only show net changes in population and miss any movement between block groups.

Columns 2, 4, and 6 of Table 2.4 reveal some of the movement between neighborhoods by allowing the Crisis to have a differential effect by the block group's educational composition. This is done by interacting the instrument with the fraction of adults in each education group, yielding 9 estimates that are akin to a transition matrix. Column 2 shows the change in low-ed residents by neighborhood type. The coefficients are all similar to one another and to the estimate in Column 1, suggesting that low-ed residents left improved neighborhoods whatever the neighborhood's original composition. On the other hand, row 1 of columns 4 and 6 show that neighborhoods with a high share of low-ed residents saw a significant increase in middle-ed residents and no significant change in high-ed residents. This is consistent with the model, where an improvement a community draws in new residents from the community just above it in the stratification.

In middle-education neighborhoods, the effects are also consistent with the model. First, as Column 6 shows, high-ed people flow into these neighborhoods as expected. The behavior of low- and middle-ed people is less clear. The effect for low-ed people in Column 2 is large, negative, and similar to Column 1, but very imprecise. This imprecision could be due to the relatively small number of low-ed people (there are roughly half has many low-ed as middle-ed) or because there is a lot of variability in the response of low-ed residents; per Proposition 2, the lower-income residents may not always flee an amenity improvement. This argument also holds for the effect on the middle-ed population in originally middle-ed communities. The corresponding coefficient in Column 4 is also negative and imprecise, but smaller in magnitude than its counterpart in Column 2, which aligns with the intuition of the model that middle-income households may not be displaced as easily as low-income ones. The imprecise and negative coefficient in Column 4 could also be due to heterogeneity of preferences, which has a significant implications for the effect on high-education neighborhoods.

The effect on high-education neighborhoods is consistent with the model, though heterogeneity of preferences must be considered. First, in Column 2, we again see an imprecise negative estimate. This is not too surprising since, as Figure 2.5 shows, very few low-ed people live in neighborhoods that are predominately high-ed. Second, in Column 4, we see middle-ed people flow into the area, as expected. Third, in Column 6, we see a large and significant flow of high-ed people out of high-ed neighborhoods.

This is easily explained by nature of the pollution shock and heterogeneity of preferences. While an ideal research design would randomly assign both positive and negative pollution changes across all neighborhood types, the shock to pollution caused by the Crisis was universally negative. This means that treatment intensity is limited by initial exposure because a neighborhood cannot remove more pollution than it has, which creates an asymmetric correlation between treatment intensity and preferences across income types. A lower-income neighborhood that is highly polluted could be feasibly populated by low-income households from across the preference distribution because income effects dominate substitution effects for these households.

On the other hand, if a higher-income household is living in a highly polluted area, it is unlikely because of income effects. Either the polluted area coincides with other amenities which are attractive to the high-income household, the household has weak preferences for clean air, or both. Thus, since a high-income household must have been living in a high-pollution area in order to be treated, treated high-income households are disproportionately more likely to have weak preferences for clean air and consequently have a low MWTP for pollution reductions. This reconciles the effects in the last row of Table 2.4.

The results for rents and house price effects by initial education, found in Table 2.5, are similar. The effect on middle-ed neighborhoods is exactly as expected, with prices and rents increasing by similar proportions. The effect on high-ed neighborhoods is negative, counter to the model and intuition, but the coefficients are very imprecise. Similarly, the house price effect, which is the more relevant metric for high-income households, is also insignificant

Table 2.5: *House Prices and Rent by Neighborhood Education*

	(1) ln Price	(2) ln Rent	(3) ln Rent
Aermod_pre×post×			
% Less than HS in 2000	-0.0065	0.0014	0.0038
	[0.0045]	[0.0063]	[0.0066]
% High School in 2000	0.0131***	0.0144*	0.0145**
	[0.0032]	[0.0075]	[0.0070]
% More than HS in 2000	-0.0054	-0.0138	-0.0130
	[0.0035]	[0.0093]	[0.0084]
Unit of Observation	House	Block Group	
R^2	0.95	0.92	0.95
N	41,771	3,162	3,162

Notes: All regressions include unit-of-observation fixed effects. Controls for house-level regression in Column 1 include year-quarter effects and quadratic time trends by local geography and year 2000 SES variables (see Section 1.4.3). Controls for block group–level regressions in Columns 2–3 are the same as in Table 2.3, plus a control for year 2000 median rent. Standard errors, clustered by tract, in brackets: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

economically. For low-ed neighborhoods, the response is mixed, with imprecise effects close to zero.

Again, a likely explanation is the other amenities available in these areas. Some low-ed areas are too dirty for higher-income residents to consider, but otherwise have a reasonable bundle of amenities. These areas will see large price increases and inflows of residents after an amenity improvement. Conversely, areas with all-around poor amenity bundles will see much smaller price effects from pollution reduction because the marginal household is still a low-income.

2.5.4 Home-ownership and Incidence

As discussed in Section 2.2.3, incumbent renters who leave their neighborhood after an amenity improvement likely experienced a welfare loss due to the improvement.[16] Because the incentive to flee is driven by income effects, conditional on preferences for the amenity, lower-income households are those that will flee.[17] And while incumbent homeowners unambiguously gain from the improvement, lower-income households are less likely to be homeowners. This makes it possible for air quality improvements to be strongly regressive, depending on the home-ownership rates of low-income households and the size of the amenity improvement. Even if few low-income families own their home, if the increase in home value is extremely large for the few that are owners, it may still result in a net aggregate gain for low-income households.

Figure 2.6a plots the results of a local linear regression of a block group's home-ownership rate in 2000 on its median household income in 2000, weighted by the block group's population in 2000. Unsurprisingly, the ownership rate is strongly correlated with

16. The model implies that these households are unambiguously worse off. However, in the presence of failures in agent optimization, this may not be the case. For example, if heads of household do not fully consider the long-term costs of pollution exposure on their children's health or human capital accumulation, these households may actually be better off if they are displaced to a cleaner area. In general, panel data on individual households is needed to draw strong conclusions on this point.

17. While we see in Table 2.4 that some high-income households also flee amenity improvements, this is due more to substitution effects than income effects, which has very different implications for welfare. Additionally, many of the high-income emigrants will be home-owners, as shown below.

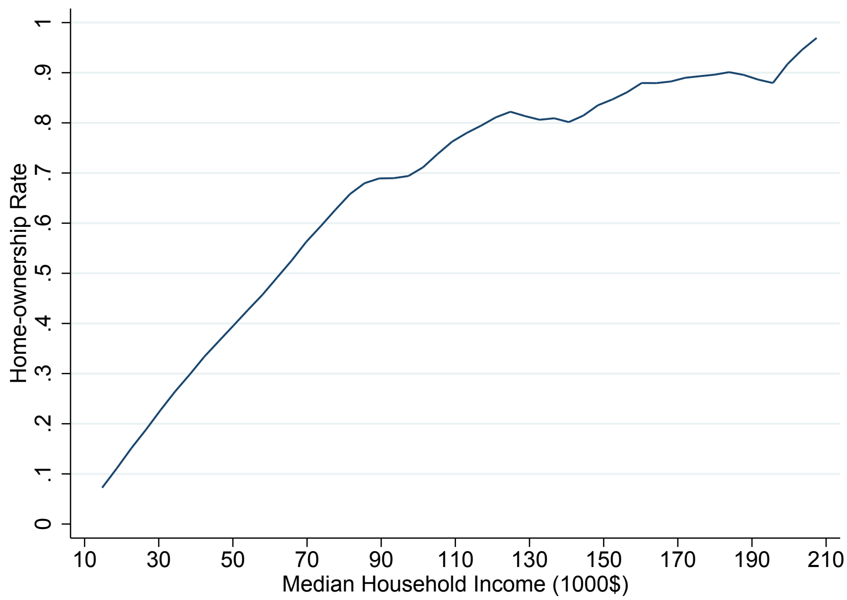

(a) *Home-ownership rate*

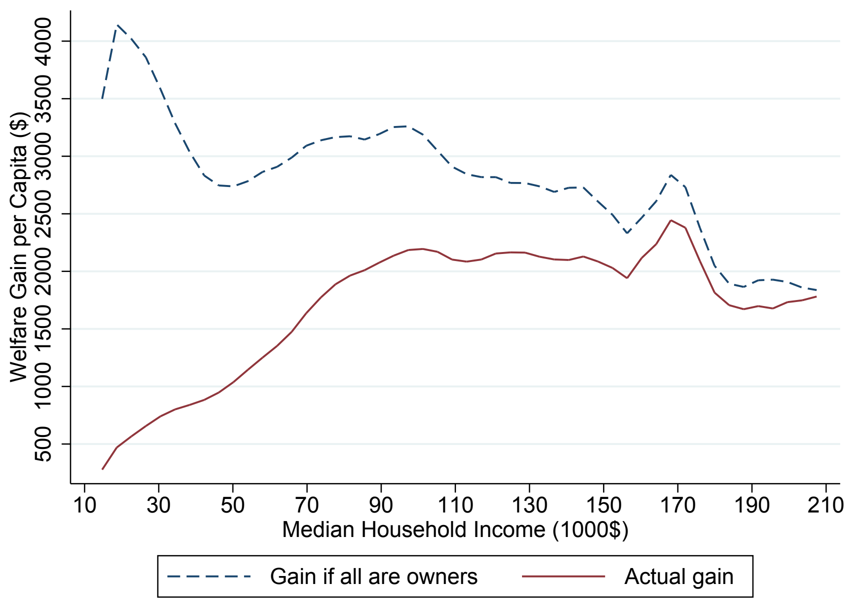

(b) *House Price Windfall*

Figure 2.6: *Home-ownership and Price Windfall by Income*

Notes: Plots are the result of local linear regressions using an Epanechnikov kernel with bandwith of 5. Sample is Census 2000 block groups, weighted by population. In subplot B, the dashed line is the gain to owners of units occupied by a household with the given income, and the solid line is the gain to residents.

income, increasing from around 10% for the poorest neighborhoods to about 95% for the richest. However, poorer areas may have also experienced larger air quality improvements since these areas were, on average, more polluted to begin with.

The dashed line in Figure 2.6b plots house-price windfall per capita by income using a local linear regression and the housing sale–based MWTP estimate from Chapter 1. This dashed line represents the per capita gain through home values if all households were incumbent homeowners. As the figure shows, poorer areas did indeed see much larger gains on average: the lowest-income areas see a gain of $3,500–4,000 per person, while the highest-income areas receive roughly $2,000.

But this differential is not enough to offset the much wider gap in home-ownership rates. The solid line in Figure 2.6b plots the windfall per capita for local owner-residents only. In the extreme case where all households are marginal and house prices capture all the welfare gains of the pollution clean up, the plot shows that the clean up is indeed regressive. In the more realistic case with many inframarginal households, it is more difficult to say. It is also possible that some non-resident landlords are themselves low-income; however, this would be a somewhat unusual scenario.

2.6 Conclusion

This paper examines how neighborhood composition and rents change in response to changes in air quality. It extends Epple, Filimon, and Romer's (1984) widely used model of spatial equilibrium to show that lower-income households may be made worse off by an improvement in their community and leave as a result. It then uses block group data from the Census to estimate how households responded to the exogenous change in pollution levels caused by the California Electricity Crisis of 2000.

The empirical results suggest that housing costs increase just as much for renters as they do for owners and that low-income households readily avoid these additional costs. Furthermore, few low-income households in the data were incumbent homeowners, making it far more likely that they were fleeing increased housing costs rather than re-optimizing after a windfall gain through home value appreciation. This provides evidence of another way air quality policies are potentially regressive (Bento 2013), and contrasts with recent work arguing that a muted rent response could make air quality policies at least somewhat progressive. More generally, the paper contributes to the extensive literature on spatial amenities, gentrification, and displacement.

Chapter 3

Lawyers Steer Clients Toward Lucrative Filings: Evidence from Consumer Bankruptcy[1]

Having a poorly informed buyer rely on the seller for market information seems like a recipe for disaster. Yet when faced with complicated decisions, many consumers rely on counsel from knowledgeable professionals who themselves have a financial stake in the buyer's decision. A growing literature documents the problems this conflict of interest can cause in myriad services, from surgeons and obstetricians to real estate agents, mutual fund managers, and funeral directors (Fuchs 1978; Gruber and Owings 1996; Levitt and Syverson 2008; Chevalier and Ellison 1997; Harrington and Krynski 2002). In each case, evidence suggested that professionals subordinated the welfare of their client to increasing their own profit, a classic example of the principal–agent problem. Here we consider the same problem for one of the preeminent advice-giving occupations—lawyers.[2] Do lawyers,

1. Co-authored with Frank McIntyre and Laura Summers.

2. While there has been some work on steerage in bankruptcy filings, discussed below, we could find little empirical work on the general steerage problem for lawyers. Ashenfelter and Bloom (1993) modeled the prisoner's dilemma problem of retaining a divorce lawyer, and Halla (2007) finds that retaining lawyers in a divorce case does little to change the final distribution of assets. Both these, and others like them, are about the

who are paid to give counsel, systematically change their counsel in ways that are likely to increase their profits? We present evidence that they do.

We consider lawyers who advise consumer debtors about whether to file bankruptcy under the cheaper Chapter 7 or the more expensive Chapter 13. The attorney's fee for Chapter 13 is set by the court rather than the market and varies greatly across bankruptcy districts, from a low of \$1500 in North Dakota to a high of over \$4400 in Nevada, despite the general uniformity of filing procedures across states (which are mostly determined at the federal level). More tellingly, the ratio of average Chapter 13 fees to average Chapter 7 fees ranges from 1.27 in the western district of Wisconsin to 5.91 in the southern district of Illinois. We show that while the optimal Chapter choice depends on the debtor's circumstances, the relative number of Chapter 13 filings is systematically *higher* in districts where attorneys are allowed to charge more for Chapter 13.

We use zip code–level data on filing rates and new data on attorney's fees taken directly from filed bankruptcy petitions to estimate this effect, which is robust to state-level fixed effects and an extensive array of zip code–level socioeconomic factors. It is also robust to controlling for the pro-creditor or pro-debtor preferences of the individual bankruptcy courts as proxied by the average repayment amount in Chapter 13 plans. If we take the estimates as causal, we find that increasing the Chapter 13 fee by 10 percent increases the fraction of bankruptcies filed under Chapter 13 filings by about 3 percent. We argue that the most plausible explanation is that some lawyers are steering debtors toward the relatively more expensive option and likely making the household worse off.[3] This conclusion, as well as the incentives that generate the steerage in the first place, is rooted in the peculiarities of

decision to retain a lawyer, rather than the quality of counsel the lawyer gives.

3. Braucher (1993), Neustadter (1986) and Sullivan et al. (1988; 1994) all provide important qualitative evidence that lawyers often steer households to file under a particular chapter of the bankruptcy code. Lefgren, McIntyre, and Miller (2010) present household-level econometric evidence from three districts taking as given the behavior of the lawyer and seek to find out how lawyers influence clients. They show that the lawyer matters heavily for what chapter gets filed, but given the available data they draw no conclusions as to why or how the lawyer does this. In this paper, in addition to bringing to bear a nationally representative data set with a cleaner identification strategy, we can explicitly identify the link between the lawyer's influence and their financial motivation.

the U.S. Bankruptcy Code itself.

3.1 Bankruptcy in the United States

Bankruptcy is the legal mechanism for dealing with insolvent debtors. Though the improving economy has seen a decline from the 2010 high in filings, almost four out of every thousand people in the United States filed for personal bankruptcy in 2012 with the average filing household holding $133,000 in assets and $206,000 in debt (Administrative Office of the US Courts, 2012).[4]

Most individual debtors file under Chapter 7 or Chapter 13 of Title 11 of the United States Code. Under both Chapters the filer must provide detailed information on all his debts, assets, income, and monthly expenses. While the bankruptcy case proceeds, creditors must cease all debt collection efforts, including foreclosure and wage garnishment. Beyond these points, the details of each Chapter differ greatly and the choice between them depends heavily on the filer's financial situation and personal preferences.

Under Chapter 7, or "liquidation", filers receive a discharge of most unsecured debts in exchange for their non-exempt assets. Asset exemptions are set by each state with the largest generally being the homestead exemption (see Li, White, and Zhu 2011). Non-exempt assets are turned over to an agent of the United States Trustee who sells them and distributes the proceeds to the debtor's unsecured creditors.[5] Once liquidation has occurred, eligible unsecured debts are discharged and all other debt contracts (including foreclosure if the debtor's home was not liquidated) proceed as usual.[6] Because debt discharge is contingent

4. Average secured debt was $126,000 while average unsecured debt was $80,000.

5. If an asset is partially exempt it is still sold by the trustee and the debtor gets the value of the exemption. For example, if the debtor owns a $100,000 house outright and the homestead exemption is $25,000, the house is sold and the debtor receives $25,000 with the remaining money going to sale costs, the trustee's fee, and unsecured creditors as usual. If the trustee deems the sale will yield very little for creditors after transaction costs and fees, he can return the asset to the debtor. If an asset is collateral for a secured debt, the sale of the asset must generate enough money to cover the balance of that debt.

6. Non-dischargeable debts are specifically named by Congress either for public policy reasons (e.g., student debt) or because the debts were incurred because of the malicious or negligent actions of the debtor (e.g., unpaid domestic support, court fines and penalties). Student debt may be discharged if the debtor can prove "undue

on asset forfeiture, Chapter 7 is generally best for people with few assets above their state's exemption limit (White and Zhu 2010). Additionally, filers who may be abusing the provisions of Chapter 7, as determined by a means test or judicial discretion, may not file under that Chapter.[7] These debtors must file under Chapter 13 if they wish to pursue bankruptcy protection.

In a Chapter 13 filing, the debtor receives a discharge of most unsecured debt in exchange for following a court-approved repayment plan.[8] In contrast with the proceeds from Chapter 7 liquidation, money from the plan goes primarily toward secured debts and may be used to pay arrears, potentially saving the debtor's home or vehicle from repossession (11 U.S.C. §1322(b)(3)). The value of many secured loans may also be reduced to the collateral's current market value.[9] These provisions present an opportunity to save one's assets, especially a home, and form the primary financial motive for filing under Chapter 13 (White and Zhu 2010). Without assets to save, a Chapter 13 often does not make much financial sense because in such cases liquidation will be far cheaper than committing to a repayment plan.

The repayment plan is drafted by the debtor and his attorney and consists of two primary elements, the number of months a plan with last and a proposed monthly payment amount.

---

hardship." See 11 U.S.C. §523 for details and the full list of non-dischargeable debts.

7. Authority for judicial discretion in preventing abuse is given by §707(b)(1). The means test was introduced by the Bankruptcy Abuse Prevention and Consumer Protection Act (BAPCPA) of 2005 and is found in §707(b)(2)(A). First, the debtor's projected five-year disposable income (projected monthly disposable income multiplied by 60) must be less than the greater of $6,575 and 25 percent of the filing's unsecured debt. Five-year income must also be less than $10,950, regardless of the level of debt (§707(b)(2)(A)(i)). For a more thorough discussion of the details and effects of BAPCPA, see Mann and Porter (2010) and Layton, McIntyre, and Sullivan (2010).

8. Discharge under Chapter 13 covers more kinds of debt than Chapter 7, including debts for willful and malicious injury and certain taxes, but only if the filer successfully completes the repayment plan. (Plan completion is sufficient but not necessary to receive a discharge, see below.) Otherwise, the list of dischargeable debts is the same as in Chapter 7. See 11 U.S.C. §1328.

9. This is known colloquially as "cram down" and can be done with any secured loan, except a first mortgage on a primary residence, where the total amount owed is greater than the market value of the asset (§1322(b)(2)). This includes secondary liens on the filer's primary residence, known as "lien stripping." However, in order to take advantage of the cram down, the new crammed-down balance of the secured debt must be paid in full by the plan. The difference between the original and crammed-down values of the debt is treated as unsecured debt. If the debtor fails to complete the repayment plan and the case is dismissed or converted, all debt that was secured when the bankruptcy petition was originally filed is again considered secured (§349(b)). In general, see §349 and §1322 or White and Zhu (2010). Cram down is also limited to assets that have been held by the filer for some minimum length of time depending on the type of asset (see §1325(a)).

In practice, most plans commit to the statutory maximum term of 5 years, though shorter terms are allowed.[10] The monthly payment amount is the debtor's calculated disposable income: predicted monthly income less "reasonably necessary" expenses for the debtor's family, charitable contributions, and business expenses (§1325(b)(2)).[11] The debtor, with the help of his attorney, calculates his disposable income, his plan length, lists how much he plans to pay to his creditors, and submits the plan to the court which ultimately decides what expenses are "reasonably necessary" and whether the various statutory requirements of the plan have been fulfilled. Of these requirements, two have particular economic significance. First, all secured and priority unsecured debts due by the end of the plan must be repaid in full. If not yet due, these debts must be kept current. Secondly, non-priority creditors must receive at least as much money as they would have under a Chapter 7 filing (§1325(a)(4)). Because of these requirements, a filer with lots of assets or secured debt but little income may not be able to afford a Chapter 13.

Even if the plan meets the all the statutory minimums, the court may still reject it. For instance, some courts require all plans be paid via wage order (similar to wage garnishment) even though the law does not require this (§1325(c); Norberg and Velkey 2006). Some courts set a higher minimum amount of unsecured debt that must be repaid.[12] Local lawyers know these minimums and craft their clients' plans accordingly (Braucher 1993). This gaming of the debtor's budget can be so severe that a plan can provide no money for living expenses (that is, 100% of gross income goes to the plan) and still be confirmed despite its apparent infeasibility (Norberg and Velkey 2006).[13]

10. The statutory minimum commitment period is 3 years, but debtors who fail the Chapter 7 means test must commit to the maximum 5 years (§1325(b)(4)(A)). If the debtor can fully repay all allowed unsecured claims more quickly than the minimum term, the minimum may be relaxed (§1325(b)(4)(B)).

11. Some types of income, such as child support, are exempt. Additionally, if the debtor fails the Chapter 7 means test, some standardized IRS expense amounts are used as a baseline for the "reasonable expenses," though additional allowances can be requested by the debtor (§1325(b)(3); §707(b)(2)(A)(ii); §707(b)(2)(B)).

12. The court is authorized to dismiss a plan made in bad faith (§1325(a)(3)). However, the legality of court-specific standards more stringent than those set forth by the Bankruptcy Code is questionable. See Norberg and Compo (2007), especially footnote 57.

13. Norberg and Compo relate the "heartbeat" test used to determine plan feasibility by one chief bankruptcy judge in their sample: "if the debtor has a heartbeat, the plan is feasible" (p. 509).

After the debtor's plan is confirmed by the court, he begins making payments. Upon fulfillment of the plan, his remaining unsecured debt is discharged and ongoing secured debt contracts (e.g., mortgages) continue according to their original terms. If the debtor fails to complete his repayment plan, the judge can grant a discharge anyway if certain conditions are met.[14] More often, the case is either converted to a Chapter 7 or dismissed altogether (§1307).

One last distinction between the two Chapters—and the motivation behind this paper—is how the debtor's attorney is paid. For a Chapter 7 filing, the attorney is generally paid upfront in cash (Gross, Notowidigdo, and Wang 2012). Chapter 7 fees are subject to court scrutiny but, as we will argue, the fees are largely market determined.[15] For a Chapter 13, the fee structure is very different. First, Chapter 13 fees are higher due to the increased complexity of Chapter 13 cases.[16] Second, the attorney receives little money upfront and is instead paid through the repayment plan as a priority creditor (§330(a)(4)(B); §507(a)(2)).[17] This is done for the benefit of filers, who are usually liquidity constrained, and to incentivize the attorney to continue putting forth effort through the much lengthier Chapter 13 process (Norberg and Compo 2007).[18]

However, because the attorney is now paid by the bankruptcy estate, her fees are subject to more stringent regulation. Attorneys must submit an itemized report of services rendered,

14. These conditions are that non-payment was due to circumstances beyond the debtor's control (e.g., job loss), unsecured creditors have received at least as much as they would have under Chapter 7, and modification of the plan to accommodate the unforeseen circumstances is not practicable. See §1328(b).

15. Because of the fragile financial state of most filers, filings under both Chapters must include a statement of compensation listing how much the attorney has been and will be paid. If the total compensation negotiated by the filer and attorney exceeds a "reasonable value" for the services rendered for the specific case, the court may alter the compensation agreement (§329).

16. In our complete sample, the average Chapter 7 fee is $1,039 while the average Chapter 13 fee is $2,513.

17. The first paid unsecured claim is trustee's administrative expenses if such payments are required to facilitate payment to other claim holders. The other "first priority" claims are domestic support obligations. Second priority payments are the attorney's fees, along with other administrative expenses of the case listed in §503(b), unsecured claims of any Federal Reserve Bank authorized under §13(3) of the Federal Reserve Act, and fees assessed under chapter 123 of Title 28 U.S.C. See 11 U.S.C. §507 for the full list of priorities.

18. Braucher (1993; 1997; 1999) discusses the effects of allowing households filing under Chapter 13 to pay their legal fees on credit. Braucher (1993) notes that in cases that last through the first payment distribution, lawyers collect at least as much as they do in a Chapter 7 case.

who within the firm performed those services, that person's standard hourly rate, etc., so the court can determine what portion of the fee is "reasonable" (§330(a)(3)). Digging into the minutia of the attorney's billings on every case is onerous, and most courts have introduced the option of a "no look" fee which attorneys can receive without submitting the full report of their services and without a hearing. And though the standard court approval remains available and is used when cases are especially taxing, the no-look fee predominates (see Price (2012) for a thorough description of no-look fees and a discussion of their questionable legality).

In conjunction with the treatment of the legal fees as priority claims, the practice of no-look fees, while understandable in view of the court's limited resources, may create perverse incentives for lawyers to inappropriately influence their clients in choosing which Chapter to use.[19]

3.2 Choosing Between Chapters

The mere existence of bankruptcy changes the incentives debtors face on a multitude of dimensions, including consumption and savings behavior, labor market supply, and state of residence (Dobbie and Song, Forthcoming). In this paper, we take their past behavior as given, including the choice to file bankruptcy, and consider only the choice between filing under Chapter 7 or Chapter 13. As such, all further discussion will be about relative decisions and incentives (i.e., Chapter 13 versus Chapter 7) faced by the court, the debtor, and the debtor's attorney.[20]

19. This possibility, the core question of our paper, has been discussed in the legal literature, though it does not seem to be of great concern (Price 2012).

20. For an examination of the decision between filing bankruptcy and informal default, see Dawsey and Ausubel (2004) and Lefgren and McIntyre (2009). Eraslan, Li, and Sarte (2007) estimate a structural model of Chapter 13 outcomes (e.g., what determines how much creditors recover of their claims). Zhu (2011) describes how a debtor's basic asset position and demographics affect various bankruptcy-related decisions, including chapter choice.

3.2.1 The court's decision

The court's stated objective is to ensure all parties have an opportunity to assert their rights under the law. The individual judge or judges in a district may also have idiosyncratic preferences about how much non-priority debt should be repaid, which will in turn affect how lenient they are approving repayment plans, how they set Chapter 13 fees, and whether they allow a debtor to file under Chapter 7 or at all.[21]

The relationship between a court's preferences for debt repayment and their leniency in approving repayment plans is straightforward. If the court thinks debts should be repaid they will require more repayment from debtors. The same holds for the court's decision to push filers wholesale from 7 into 13. Thus, net of the indirect effects on the debtor's decision, we should see a positive correlation between how much of the debtor's income is earmarked for the repayment plan (which we will call the "repayment rate") and the fraction of bankruptcies filed under Chapter 13.

How pro-creditor preferences determine the no-look fee is less clear. On one hand, the court may set the fee low in order to entice debtors into Chapter 13 or to directly transfer money from attorneys to creditors. On the other hand, the court may feel creditors are best served if the debtor has good (or well motivated) legal representation and therefore set a high fee.[22] Because of this ambiguity, we explore this issue further in section 3.4.2.

3.2.2 The debtor's decision

The debtor faces a discrete choice that is determined by his idiosyncratic tastes for each chapter, his financial situation, the price of filing under each chapter, and the advice he gets from his attorney.

Idiosyncratic taste for Chapter 13 could include a moral conviction to repay one's debts

21. Sullivan, Warren, and Westbrook (1994), Norberg and Velkey (2006), Norberg and Compo (2007), and Dobbie and Song (Forthcoming) discuss the geographic variation in culture and practice of bankruptcy law.

22. Braucher (2001) finds that higher attorney's fees were correlated with plan completion, though Norberg and Velkey (2006) find no relationship. Regardless of the true empirical relationship, it is not hard to imagine legal actors believing the former and acting accordingly.

or non-monetary desire to save one's house from foreclosure (e.g., the house has sentimental value, desire to avoid moving school-age children). These preferences are unlikely to be captured in most data, including ours. Nevertheless, in section 3.4.2 we consider several potential proxy variables for these unobserved preferences to measure their importance relative to the core financial aspects of the case.

Calculating the financial benefits of filing under Chapter 13 is more or less an accounting exercise. The debtor retains his non-exempt assets in exchange for his multi-year repayment plan, which may include paying arrears and avoiding repossession.[23] These payments are his income minus court-approved expenses. As previously mentioned, this trade-off between liquidation and the repayment plan only makes financial sense if the ratio of assets to repayment is high. If the ratio is too low, the amount the debtor pays to non-priority creditors beyond what he would have under the Chapter 7 will be more than his assets are worth.

As this asset-to-repayment ratio increases, more of the debtor's repayments go toward keeping his own assets rather than to non-priority creditors whose claims would be discharged under Chapter 7 anyway. On the other hand, with too many assets the debtor would not be able to meet the statutory requirement of Chapter 13 to cover both his priority debts and his non-priority debts' Chapter 7 counterfactual payments. These conflicting effects create a range for the total repayment amount over which the debtor, conditional on his idiosyncratic preferences, prefers Chapter 13.

Because this optimal Chapter 13 range is continuous while the Chapter decision is discrete and the repayment amount is relatively fixed for a given filing, the debtor should be fairly insensitive to the relative prices since Chapter 13 fees are paid through the plan. This follows from the fact that for a fixed repayment plan amount, and in the typical case where non-priority debts are only partially repaid, any increase in legal fees is a one-to-one transfer from non-priority creditors. Therefore, the debtor is indifferent to these fees as far

23. There are also more kinds of debt eligible for discharge under Chapter 13 and the cram down provisions of Chapter 13 could be very appealing to some debtors. However, as the expected value of these provisions are small relative to secured and consumer debt, we do not consider them here.

as he is indifferent between the welfare of his creditors and his lawyer.

Of course, the debtor may not be completely insensitive to the relative price of a Chapter 13. Chapter 7 fees are paid upfront in cash and are fully borne by the debtor. Also, if a Chapter 13 case is dismissed the debtor forfeits any legal fees that have already been paid through the plan, funds which would have otherwise paid down debt. Thus, the debtor's sensitivity to the price of a Chapter 13 is proportional to the subjective probability he assigns to dismissal multiplied by the rate at which he expects to repay his non-priority debts after dismissal. As less than half of first-time Chapter 13 filers actually receive a discharge, a fully rational debtor should see dismissal as a real possibility (Dobbie and Song, Forthcoming). More realistically, the debtor will be subject to any of myriad cognitive biases, such as optimism bias, that push his perceived probability of dismissal closer to zero (Van den Steen 2004; Brunnermeier and Parker 2005; Mayraz 2013). His perceived probability could also be influenced directly by the advice of his attorney.

The final factor that shapes the debtor's Chapter choice is the counsel proffered by his attorney. The debtor likely depends heavily on this counsel when calculating the financial details previously discussed or when choosing a Chapter in general. Whatever the exact mechanism, the attorney has a strong hand in determining Chapter choice (Lefgren, McIntyre, and Miller 2010).

3.2.3 The attorney's decision

When choosing how to advise her client, the attorney maximizes some balance of her own profits from the case and her client's well-being. These preferences could be complex (e.g., altruistically lexicographic in the client's well-being) or mundane (e.g., zero weight on either profits or client). If Chapter 13 prices were set in a purely competitive market, the attorney's economic rents for both Chapters would be zero. Since she is compensated for her time at the market rate regardless of Chapter choice, she would be financially indifferent between the two. Therefore, we should only see undue steerage by attorneys if the price is not set by the market and the attorney puts positive weight on her own profits/rents.

Even under these circumstances it may be true that because of the high probability of plan dismissal, attorneys are themselves somewhat insensitive to the price, or they require large nominal fees so that in expectation they receive a reasonable wage. This is unlikely given the high priority attorneys have among claimholders. Because of this high priority, attorneys still recover a large portion of their fee even when the filer fails in his repayments. And since they are paid in full when a repayment plan succeeds, it appears that the expected value of a Chapter 13 fee is very close to the nominal fee.[24] Thus, if Chapter 13 fees are set in a noncompetitive fashion and attorneys are concerned about their own profits, an increase in the court-allowed Chapter 13 fee should encourage the attorneys to exert pressure on debtors to file under Chapter 13.[25]

3.2.4 Summary of predictions

In equilibrium, the positive price response from attorneys will be countered by the (potentially small) negative price response of filers. From the arguments presented above, we make the following predictions: Suppose Chapter 13 fees are not set by the market. Then, if attorneys try to push clients to serve their own interests and debtors are insensitive to the relative price of a Chapter 13, we should see a strong positive relationship between the relative quantity of Chapter 13 filings and their relative price. If attorneys do not try to push clients toward the more lucrative filing, we should see a non-positive relationship between relative quantity of Chapter 13 filings and their relative price.

While a positive relationship between relative quantity and relative price of Chapter 13

24. For example, Lefgren, McIntyre, and Miller (2010) look at 54 dismissed Chapter 13 filings drawn from the northern bankruptcy district of Texas. The average case was dismissed within 13 months, yet lawyers collected 66% of their nominal fee through the few repayments that were made. In 22 of the 54 dismissed cases, the lawyer was paid more than all other creditors, secured and unsecured, combined. Braucher (1993) also finds evidence of heavy front-loading (see footnote 18 in this paper). Using our data, we spot checked filings in eight bankruptcy districts and found that in five cases, lawyers were paid prior to unsecured creditors. In two cases ordering was unclear and in only one case was the lawyer clearly paid pro rata throughout the Chapter 13 plan.

25. It may also be the case that attorneys are more likely, whether for practical reasons or reasons of conscience, to be more responsive to price incentives when the debtor's financial situation is more suited to Chapter 13. That is, there is an interaction term in the attorney's objective function between the price incentive and the client's interest. We explore this possibility further in Section 3.4.3.

is evidence of steerage, it does not immediately imply a principal–agent problem. A market with a binding price ceiling would also see the quantity supplied increase when the ceiling is raised. Appendix C.1 considers the substantial empirical problems with that approach as opposed to a principal–agent problem.

Even if lawyer steerage is the result of a principal–agent problem and not a binding price ceiling, our predictions above still depend on the supposition that judicial price setting and debtors' price insensitivity keeps Chapter 13 fees from being generated by the market. One sign of this non-market price setting, discussed more in the next section, is the far lower within-district variability in Chapter 13 fees compared with both Chapter 7 and the large variance across districts. Chapter 13 filings are not only more complex than Chapter 7s, they are also more idiosyncratic and variable. Just as the complexity of an individual's tax return scales with the complexity of their financial situation, so too does the complexity of a Chapter 13 filing. We might expect a simple Chapter 13—a debtor with only a mortgage in arrears and some credit card debt—to be easier and cheaper than one with multiple properties, a small business, secondary liens that may be crammed down or stripped, etc. At the very least, the complexity of Chapter 13 filings is no less varied than that of Chapter 7 filings. Yet, as we discuss below, this complexity does not convert into higher price dispersion across individuals, suggesting that normal market forces are not at work.

3.3 Empirical Strategy

3.3.1 Specification

The ideal empirical strategy would mirror the debtor's discrete choice and estimate the parameters of that model. Unfortunately, a lack of individual-level data prevent this. We can approximate the individual's discrete decision by looking at the fraction of bankruptcy filers in a zip code who choose Chapter 13. Consider the following regression equation:

$$\ln\left(\frac{f_i^{13}}{f_i^T}\right) = \delta p_d + \beta_1 r_d + \beta_2 i_d + \beta_3 S_d + \Gamma_1 X_i + \Gamma_2 W_s + \varepsilon_i \tag{3.1}$$

where i is the zip code, d is district, and s is state. f_i^T is the number of total bankruptcies, f_i^{13} are those filed under Chapter 13. p_d is the log of the district's ratio of prices for filing a Chapter 13 and a Chapter 7 in 2007, $\ln\left(p_d^{13}/p_d^7\right)$.[26] The remaining covariates are designed to control for differences in socioeconomics and local legal practices.

The variable r_d is the district's average log repayment rate calculated from our sample of Chapter 13 filings from PACER.[27] One may be concerned that the level of Chapter 13 fees is correlated with the court's unobserved policies favoring or disfavoring creditors. In that case, p_d and ε_i would be correlated and though the estimate for δ could still be taken as a sign of the importance of legal institutions in determining which Chapter is filed, it would not estimate the price elasticity. Therefore, we include r_d, the log repayment rate, to proxy for these unobserved preferences.

It may be that this repayment variable is simply reporting a correlation between higher incomes and more Chapter 13 filings. We account for this in two ways. First, we include an extensive set of zip-level demographic variables, X_i, to control for the distribution of household income and other demographics.[28] Second, we include i_d, the district's average reported log income from our sample of Chapter 13 filings.

Because Chapter 13 is designed to save assets, most notably the debtor's home, from liquidation, we include zip code–level data for home ownership and housing prices in X_i. We also include S_d, a district average for the log fraction of assets that are secured.

Lastly, we add a vector of state-level variables, W_s, that control for the state's credit

26. Although we present OLS results here, we also tried Poisson and Negative Binomial specifications, which would better fit the data generating process and allow us to include observations with zero filings. The elasticities were identical to those reported here for the OLS estimation, with similar standard errors, and so we omitted them for the sake of simplicity. In unreported work we also ran regressions on the level of fraction 13, rather than the log value. The results were substantively similar to those presented here and omitting zip codes with zero Chapter 13 filings made no difference on the point estimates. Thus we do not believe that our specification creates any substantial bias due to it requiring us to exclude zero values.

27. All of the district-level variables are constructed using averages of PACER filings.

28. The zip-level variables included in X_i are reported in Table C.1 which gives the complete regression results for our baseline specification. These variables are urban fraction, population size, marital status variables, household composition, unemployment, self-employment and home ownership rates, housing values, and the distributions of education, race, age, and income. These either directly control or proxy for differences across areas in asset and debt levels, income, and the unobserved costs of filing bankruptcy, whether financial, social, or preference based.

and bankruptcy policies. First, we use two dummy variables for the how stringent the state's wage garnishment laws are; Lefgren and McIntyre (2009) show these laws to be an important determinant of overall bankruptcy levels. The omitted category is states that follow the federal law, which requires wage garnishments to be less than 25 percent of the worker's paycheck. The two included dummy variables capture states that impose further restrictions on wage garnishing and states that essentially ban the practice. Second, we include two dummies for varying levels of homestead exemptions as higher exemptions may make Chapter 7 more attractive. Finally, we include a dummy variable for whether the state requires judicial foreclosure proceedings.[29] In states with judicial foreclosures, foreclosure is more costly for the lien holder and is therefore less of an immediate threat to the debtor. We would expect debtors in these states to find Chapter 13 bankruptcy less appealing.

Given the constraints of the available aggregate data, Equation (3.1) accounts for all the major components of the bankruptcy process presented in the prior section and controls for major sources of potential bias. To account for state- or district-level correlation, we always cluster the standard errors at the state level. The regressions are weighted by the number of bankruptcies in the zip code in order to be representative of the bankrupt population.

Unfortunately, W_s cannot perfectly control for the full range of state policies or preferences that may affect bankruptcy choices. We thus also employ a more stringent identification strategy that replaces the state controls with state fixed effects. The state fixed effects should capture not only political economy variables, but any differences in the economic environment across states. Additionally, if state-level demographics affect bankruptcy policy through political economy channels, this will also be captured by the fixed effects. With state fixed effects, Equation (3.1) becomes

$$\ln\left(\frac{f_i^{13}}{f_i^T}\right) = \delta p_d + \beta_1 r_d + \beta_2 i_d + \beta_3 S_d + \Gamma_1 X_i + \mu_s + \varepsilon_i \tag{3.2}$$

29. In a state that requires judicial foreclosure, a lien holder must formally file suit against the delinquent borrower and appear in court before eviction and sale can take place. This additional cost cuts the probability of foreclosure roughly in half (Mian, Sufi, and Trebbi 2014).

In this case, our identification of δ shifts from comparing districts across states, controlling for differences in state laws, to comparing districts within a given state. We still control for zip code–level differences and observable differences in district policy. In this specification, a state with only one district will have no within-state variation in the district variables such as Chapter 7 or 13 fees. Those single-district states will help identify the effects of zip-level covariates only. Thus, for Equation (3.2) our identification of district-level effects will rely on the 24 states with more than one district.[30]

3.3.2 Data

Our main source of data is zip code–level bankruptcy counts obtained from Lundquist Consulting. Zip code demographic information was taken from Census data. We match each zip code with the bankruptcy district where its residents file. When a zip code contains filings in multiple districts, we match it to the district where the most bankruptcies from that zip code were filed.[31]

These zip code–level counts do not include any other information about the filers or their financial situation, so we also match the counts from 2007 with the district-level averages of legal fees, repayment rates, filer income, and fraction secured debt taken from a random national sample of approximately 1,500 bankruptcy cases filed the first week of February 2007 obtained through the U.S. court database, PACER.[32] Using 2007 data puts us more than a year after the 2005 legal reform, thus allowing time for adjustment to the new regime. We supplemented this original source with additional filings drawn directly from PACER to ensure we had at least two Chapter 7 filings and four Chapter 13 filings, making sure each listed a non-zero fee paid to the lawyer. The repayment plan is not available in the initial

30. Table 3.6 reports on a regression where only these 24 multi-district states are included; the δ coefficient is unaffected.

31. Of the zip codes in our sample, 84 percent have no filings in a second district and 95 percent have less than 10 percent of their filings in a second district. When weighting by number of filings, 98 percent of our zip codes have less than 10 percent of filings in a second district.

32. Thanks to Michelle Miller who provided these PACER filings.

bankruptcy petition, so we use the difference between the reported current monthly income and monthly expenditures listed which, when given, is a good proxy for the plan's monthly installments.[33] The legal fees, repayment rates, filer income, and fraction secured debt were all extracted from individual PACER filings and then averaged over each district.

For state-level controls we use data on state garnishment restrictions, married homestead exemption levels, and judicial foreclosure regulations.[34] Table C.1 of the appendix lists these variables for each state, along with the state's average number of Chapter 13 filings per district that we use to estimate prices.

Table 3.1 shows summary statistics of the zip code–level bankruptcy and demographic information. We see that, on average for a given zip code, 3.02 households per thousand filed for personal bankruptcy. Of this number, 1.91 per thousand are Chapter 7 and 1.11 are Chapter 13. In the average filer's zip code of residence, 61 percent of the population are married, 71 percent graduated from high school, 4 percent are unemployed, and 55 percent are homeowners (based on 2000 census data).

Table 3.2 reports data for states with two or more districts, as these are the districts relevant to our fixed effects identification strategy. The table lists the average fraction of bankruptcies that are Chapter 13 in the state, as well as the spread between the highest and lowest averages across the bankruptcy districts in the state. Table 3.2 also reports on

33. As noted, we made sure to have at least six observations on legal fees (two on Chapter 7 and four on Chapter 13), and four on each of the other district averages used as controls. The additional filings were also pulled from early February 2007. Given the limited data, the estimates likely contain some measurement error, which we address in more detail in section 3.4.2. We also checked how well our district filings matched the census demographic data. Specifically, districts with high home ownership do show more filings with secured debts ($\rho = 0.10$) and as census median income rose, so too does the average income reported in bankruptcy filings ($\rho = 0.19$ in log income). Similarly, filing income was positively correlated with the fraction of high income households in the census and negatively correlated with the fraction of low income households.

34. We obtained wage garnishment information primarily from Fair Debt Collection.com (http://www.fair-debt-collection.com/state-wage-garnishments.html), double checking the information with The Commercial Bar (http://www.commercialbar.com/sumcoltn.htm) and BCS Alliance (http://www.bcsalliance.com/y_debt_statelaws_garnishments.html) to ensure accuracy. The laws we refer to are for wage garnishment of non-priority, private debts. Many states that we code as restricting wage garnishment still allow it in the case of, for example, child support payments. Andreas Lehnert (see Lehnert and Maki 2002) kindly provided information on 2000 exemption levels. Judicial foreclosure data are taken from Dobbie and Song (Forthcoming).

Table 3.1: *Sample Summary Statistics*

Variable	Mean	Std. Dev.
Bankruptcy Filings per Thousand People	3.02	4.01
Chapter 7 Filings per Thousand People	1.91	2.45
Chapter 13 Filings per Thousand People	1.11	2.06
Urban	0.78	0.32
Population	11006	13777
Married	0.61	0.10
Divorced	0.11	0.03
Household of 2	0.32	0.05
Household of 3	0.17	0.03
Household of 4	0.15	0.04
Household of 5	0.07	0.02
Household of 6	0.03	0.02
Household over 6	0.02	0.02
Female Head of Household	0.19	0.10
Head of Household Below Age 24	0.04	0.02
Finished High School	0.71	0.11
Finished college	0.19	0.11
Black	0.16	0.23
Hispanic	0.09	0.16
Other Race	0.05	0.07
Age Below 6	0.08	0.02
Age 6 to 18	0.19	0.03
Age 19 to 24	0.08	0.04
Age 25 to 29	0.07	0.02
Age 30 to 39	0.16	0.03
Age 40 to 49	0.15	0.02
Age 50 to 59	0.11	0.02
Unemployed	0.04	0.02

Table 3.1: *(continued)*

Self-Employed	0.11	0.04
Household Income under $10,000	0.09	0.06
Household Income $10-$20,000	0.13	0.05
Household Income $20-$30,000	0.13	0.04
Household Income $30-$40,000	0.13	0.03
Household Income $40-$50,000	0.11	0.02
Household Income $50-$60,000	0.09	0.02
Household Income $60-$75,000	0.11	0.03
Household Income $75-$100,000	0.10	0.05
Fraction Homeowners	0.55	0.16
25th Percentile of Log Housing Value	11.3	0.51
75th Percentile of Log Housing Value	11.9	0.462
Observations	25253	

Notes: Bankruptcy rates are weighted by population. Demographics (except population) are weighted by the number of bankruptcy filings.

Table 3.2: *Summary Statistics of States with More than One District*

State	Bankruptcy Rate	Chapter 13 Fraction		Chapter 13 Fee		Log 13/7 Fee Difference	
		Average	Max-Min	Average	Max-Min	Average	Max-Min
4 Districts							
California	2.05	0.29	0.20	$3,469	$1,113	1.17	0.72
New York	2.17	0.27	0.08	$2,910	$1,962	0.88	0.42
Texas	1.90	0.56	0.11	$3,041	$609	0.51	0.58
3 Districts							
Alabama	4.93	0.70	0.15	$2,571	$663	1.06	0.16
Florida	2.07	0.33	0.14	$2,926	$222	0.72	0.07
Georgia	5.64	0.67	0.27	$2,837	$1,152	1.07	0.57
Illinois	3.36	0.38	0.26	$2,890	$1,127	1.49	0.86
Louisiana	2.74	0.61	0.21	$2,225	$775	0.84	0.78
North Carolina	2.22	0.55	0.19	$2,806	$500	0.86	0.50
Oklahoma	2.37	0.21	0.13	$2,385	$1,031	0.82	0.38
Pennsylvania	2.34	0.39	0.18	$2,684	$1,381	0.95	0.34
Tennessee	6.85	0.60	0.31	$2,185	$280	0.89	0.20
2 Districts							
Arkansas	3.99	0.50	0.07	$2,550	$0	1.24	0.06
Indiana	4.76	0.31	0.06	$2,933	$633	1.00	0.44
Iowa	2.30	0.10	0.07	$2,035	$371	0.83	0.05
Kentucky	3.96	0.31	0.02	$1,963	$275	0.72	0.01
Michigan	4.03	0.25	0.15	$2,265	$1,120	0.99	0.53
Mississippi	3.76	0.58	0.02	$2,225	$50	0.87	0.33
Missouri	3.60	0.35	0.08	$2,480	$1,040	1.11	0.58
Ohio	4.30	0.34	0.06	$2,332	$879	0.86	0.54

Table 3.2: *(continued)*

Virginia	2.63	0.40	0.07	$2,834	$333	1.03	0.22
Washington	2.66	0.28	0.07	$2,071	$143	0.90	0.08
West Virginia	2.15	0.11	0.02	$2,328	$1,344	0.81	0.64
Wisconsin	2.76	0.23	0.13	$1,950	$800	0.58	0.58
All Multi-District States	3.31	0.39	0.13	$2,537	$742	0.92	0.40
All States	3.03	0.34	-	$2,513	-	0.89	-

Notes: Columns report unweighted averages across districts within the state. "Max-Min" is the within-state difference between the districts with the highest and lowest values. Averages over all states are weighted by district, not population or bankruptcy filings.

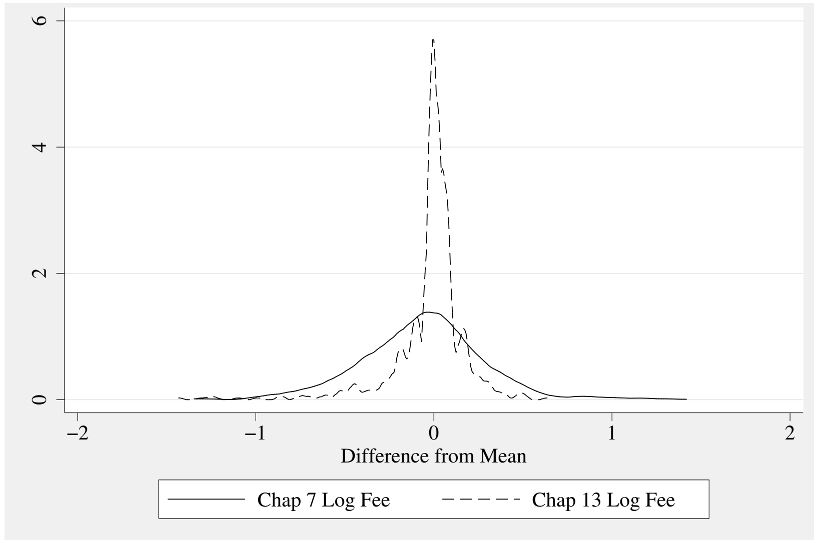

Figure 3.1: *Kernel Density of Log Fees About Within-district Mean*

filing fees, the log difference between Chapter 13 and Chapter 7 fees, and the in-state spread between districts.

In these multi-district states, the average Chapter 13 filing costs $2,537, which averages 0.92 log points more than a Chapter 7 filing. Chapter 13 rates also differ a great deal within states—typically being $742 lower in the lowest district than in the highest within a given state. One can also see that 12 states have at least three districts and about half of all states have at least two districts. The last row of the table reports comparable information for all states, where the average Chapter 13 fee is $2,513, almost identical to the multi-district average.

The distribution of Chapter 13 fees can give some indication if lawyers are in fact clumping at each court's suggested Chapter 13 fee. Figure 3.1 shows the kernel densities of log Chapter 13 fees and log Chapter 7 fees taken from our sample of filings, centered about the respective district's mean. The Chapter 13 fees in each district appear to be clumped far more tightly around the mean than the Chapter 7 fees; 65 percent of Chapter 13 fees lie within 10 percent of the district's mean, while 29 percent of Chapter 7 fees are within 10 percent of the mean. Given the higher complexity, and hence variability, of Chapter 13 filings, if we take the Chapter 7 fees as a lower bound of normal bankruptcy-market price dispersion, the noticeable clumping in the Chapter 13 fees suggests the court price substantially influences the lawyer's fees for Chapter 13 bankruptcy filings.[35]

3.4 Results

Table 3.3 reports on our specifications outlined in equations (3.1) and (3.2). Column 1 reports on a simple specification that does not control for any state-level differences. Here we see no significant relationship between fees and filing ratios. This changes dramatically once we account for cross-state variation in the next two columns.

Column 2, which uses the state-level controls discussed in Equation (1), reports the price elasticity estimate as 0.17 and is statistically significant, indicating that higher fees lead to a higher fraction of Chapter 13 bankruptcies. The effect of log repayment rate is positive, though small and not significant.[36] Note that in the debtor's optimization problem, higher repayment rates should drive down Chapter 13 filings. The positive coefficient suggests that this effect is being overwhelmed by an omitted variable. The best candidate is that repayment rates are largely being determined by the court. Courts that require more repayment are manifesting a pro-creditor culture and so are likely more inclined to push debtors toward Chapter 13. Thus the positive coefficient is a sign that repayment rates are proxying for the court's ability to affect the explicit and implicit costs of bankruptcy, in line with our prediction in section 3.2.1.[37]

Secured debt as a fraction of total debt is an insignificant predictor of chapter choice. In

35. Lawyers can petition the court's typical fee for unusually time-consuming filings, and of course they can always charge less than the maximum fee, thus we do see some variation in the reported fee. Also, Chapter 7 filings may also include some business debts, the added complexity of which would contribute to the upper tail of the Chapter 7 distribution.

36. Reported income is negatively related with Chapter 13 filings—a peculiar outcome considering Chapter 13 is typically associated with higher incomes. Of course, our unreported zip code controls include extensive household income controls, so the filing income variable is difficult to interpret by itself. Further, this result is insignificant in our later, preferred specifications. As we show in our robustness checks, the variable is also largely irrelevant to estimating the price elasticity.

37. Although we do not report here on the dozens of coefficients in the demographic controls, Table C.2 of the appendix reports on these for the preferred specification in Column 3.

Table 3.3: *The Effect of Prices on the Log Fraction of Filings that are Chapter 13*

	(1)	(2)	(3)
Log 13/7 Fee Difference	0.05	0.17**	0.32***
	[0.12]	[0.08]	[0.08]
Log Repayment Rate	0.09	0.12	0.33**
	[0.29]	[0.21]	[0.16]
Log Monthly Income	-0.26*	-0.22*	-0.17
	[0.13]	[0.14]	[0.12]
Log(Secured Debt/Total Debt)	0.15	-0.09	-0.07
	[0.13]	[0.14]	[0.22]
Judicial Foreclosure		0.00	
		[0.08]	
Some Restrictions on Wage Garnishment		-0.24***	
		[0.08]	
Heavy Restrictions on Wage Garnishment		0.26***	
		[0.09]	
Medium Homestead Exemptions		-0.10	
		[0.10]	
High or Unlimited Homestead Exemptions		-0.07	
		[0.10]	
State Fixed Effects			yes
R-squared	0.39	0.44	0.57

Notes: Dependent variable is Log Fraction 13. Each zip code in the regressions is weighted by the number of bankruptcies that occurred there in 2007. Demographic controls at the zip-code level are included in each regression and are those summarized in Table 1: urban, population and population squared, marital status, household size, gender and age of head of household, education, race, age, employment status, self-employed, home ownership, and household income. N = 20051. Standard errors in brackets are clustered at the state level. * significant at 10%, ** significant at 5%, *** significant at 1%.

fact, the point estimate is leaning the wrong way, which would counterintuitively suggest that districts where debtors have more secured assets actually tend to file relatively fewer Chapter 13 bankruptcies. State judicial foreclosure laws have a point estimate of zero. Our results are based on data in 2007, before the housing market crash was fully underway. Thus judicial foreclosure may have played a stronger role in later years, but at this point it seems to have had little direct impact on chapter choice. Thus these two measures of secured asset risk, obvious determinants of the relative benefits of the two Chapters, perform far worse at explaining Chapter choice than the amount the lawyer was paid.

The restrictions on wage garnishment should lower the felt burden of debt, as creditors have fewer options for demanding repayment, and so decrease the desire to file bankruptcy, but whether this has more of an effect on Chapter 13 or 7 filings depends on the marginal density of the unobserved preferences. Given the ambiguity of these variables' effect on chapter choice in our model, it is perhaps not surprising that the effect switches signs moving from medium to heavy restrictions on wage garnishment. Finally, the small effect of the homestead exemption is in line with the fact that few filers have sufficient assets to make the exemption differences relevant.

The state fixed effects estimate, based on Equation (3.2), in Column 3 is 0.32—higher than what we found before but only marginally statistically different (p-value=0.09)—and indicates that a 10% increase in the price gap would lead to a 3.2% increase in the fraction of Chapter 13 filings. Also, in this regression the repayment rate estimate is 0.33 and statistically significant, implying districts that require a lot of repayment in Chapter 13 also tend to have many Chapter 13 filers, again because of the local legal culture. Our monthly income and secured debt controls are again statistically insignificant. Thus our best estimate, which controls for unobserved differences across states by using only within-state variation, finds that price has a substantial effect on filing behavior.

3.4.1 Raw Filing Rates

There are two main mechanisms by which price can affect the fraction of filings that are Chapter 13: it can move agents from 7 to 13 or discourage debtors from filing at all. We test these two channels by re-estimating the fixed effects regression of Equation (3.2) with $\ln\left(f_i^j / P_i\right)$ as the dependent variable, where P_i is the population size for the zip code and j indexes either total filings, Chapter 7 filings, or Chapter 13 filings. If price, repayment rates, or state policies cause new entrants to bankruptcy, this will cause changes in total filings. The Chapter 7 regression catches movements out of Chapter 7 and potentially into Chapter 13. The total filings number is an agglomeration of the change in the two filing rates. Regressions are weighted by the population in the zip code, so that the results are representative of the U.S. population. Table 3.4 reports these results. In all cases, the regressions include our usual demographic controls and state fixed effects.

The point estimates confirm that Chapter 13 filings do go up but the effect on Chapter 7 and total filings is uncertain (though the point estimates are negative and positive, respectively). The large standard errors make it impossible to make any definitive declarations about whether higher Chapter 13 filings are from those who would not file or from those who would have filed a Chapter 7.

If the repayment rate is acting as a proxy for pro-creditor court institutions, we would expect a shift out of Chapter 7 filings and into Chapter 13 filings, with possibly fewer overall filings. The coefficients on repayment rates bear out this story but the standard errors in all three regressions are far too large to make the claim with any confidence.

3.4.2 Robustness

The above results suggest that raising the relative price of a Chapter 13 bankruptcy 10 percent would lead to a roughly 3 percent increase in the relative number of Chapter 13 filings. This section considers a number of changes to the specification of the variables, the regression equation, and the methodology to show that this result is robust to a wide range of different estimation assumptions.

Table 3.4: *State Fixed Effects Estimates of the Effect of Prices on the Rate of Bankruptcy Filings, by Chapter*

	(1) Log Rate of Chapter 7	(2) Log Rate of Chapter 13	(3) Log Bankruptcy Rate
Log 13/7 Fee Difference	-0.05	0.39***	0.08
	[0.12]	[0.13]	[0.11]
Log Repayment Rate	-0.21	0.34*	-0.04
	[0.23]	[0.19]	[0.19]
Log Monthly Income	0.14	-0.30*	-0.04
	[0.15]	[0.16]	[0.11]
Log(Secured Debt/Total Debt)	0.18	0.39	0.22
	[0.16]	[0.31]	[0.16]
Observations	24,214	20,051	25,253
R-squared	0.57	0.61	0.60

Notes: Each zip code in the regressions is weighted by its population in 2007. All regressions include state fixed effects. Demographic controls at the zip-code level are included in each regression and are those summarized in Table 1: urban, population and population squared, marital status, household size, gender and age of head of household, education, race, age, employment status, self-employed, home ownership, and household income. Sample Size fluctuations are from some zip codes having no bankruptcies of the reported type during 2007. Standard errors in brackets are clustered at the state level. * significant at 10%, ** significant at 5%, *** significant at 1%.

Table 3.5: *Alternate Fee Specifications*

	Baseline	Max 13 Fee	90th Percentile 13 Fee	13 and 7 Fees Separately	Only 13 Fee	Instrument Fee Difference with 2011 No-look Fee
	(1)	(2)	(3)	(4)	(5)	(6)
Log 13/7 Fee Difference	0.32***	0.23**	0.27***			0.46***
	[0.08]	[0.10]	[0.09]			[0.05]
Log Chapter 13 Fee				0.29**	0.29*	
				[0.13]	[0.15]	
Log Chapter 7 Fee				-0.34**		
				[0.17]		
Observations						18,094
R-squared						0.57

Notes: Each zip code in the regressions is weighted by the number of bankruptcies that occurred there in 2007. All controls are the same as the baseline regression in Table 3, Column 2. Column 6 instruments Log 13/7 Fee Difference with each district's stated no-look fee policies in 2011 as reported in Price (2012). The excluded instrument gives an F-stat of 35.9 in the first stage; the overall R-squared in the first stage is 0.71. Demographic controls at the zip-code level are included in each regression and are those summarized in Table 1: urban, population and population squared, marital status, household size, gender and age of head of household, education, race, age, employment status, self-employed, home ownership, and household income. State fixed effects are included in each regression. Standard errors in brackets are clustered at the state level. * significant at 10%, ** significant at 5%, *** significant at 1%.

Table 3.5 considers different ways that we might have specified the fee variable. Column 1 simply repeats the preferred specification from Table 3.3, Column 3. In that baseline specification we used the log difference of the average reported Chapter 13 and Chapter 7 fees. As seen in Figure 3.1, Chapter 13 fees are heavily clustered around the district mean, lending credence to the institutional view that courts set customary fees from which lawyers rarely deviate. On the other hand, if the Chapter 13 "no-look" fee is really functioning more like a maximum fee, then perhaps the correct empirical object to use is not the average but the maximum observed Chapter 13 fee in each district. Columns 2 and 3 consider this by replacing the average Chapter 13 fee with either the maximum Chapter 13 fee observed in the sample or the 90th percentile fee (to reduce exposure to noisy outliers). This causes the point estimates to decrease slightly to 0.23 and 0.27, respectively, but one can still reject a zero coefficient and one cannot reject that the estimates are the same as the baseline specification. A slight decrease is also what we would expect to happen mechanically when using variables with more variance (either from measurement error or the definition of extreme order statistics). Given the similarity of the results, and since upward deviations from the standard fee are permitted by the courts, we return to the average fee as our preferred measure.

Column 4 shows that our results are also robust to including the fees as two independent regressors. Indeed we get equal and opposite coefficients on the two fee variables, suggesting that the baseline differenced version is not doing violence to the underlying data generating process.[38] Column 5 uses the log Chapter 13 fee as the sole regressor and returns a statistically significant point estimate of 0.29, virtually identical to the baseline.

One concern in the baseline specification is that the price variable is formed by averaging observations within districts, and in many cases the number of observations averaged is small. Measurement error in the fees variable could be biasing our coefficients toward zero. This is especially true in the presence of state fixed effects, where the fixed effects may

38. This evidence of substitution between the two chapters also suggests that the results are not driven by a capacity constraint story, since debtors appear to be moving between chapters as prices fluctuate, rather than moving in and out of filing.

decrease the signal to noise ratio. We can deal with this by instrumenting the covariate with a second measure, even if the second measure is also measured with error. If the measurement errors in the two variables are uncorrelated, the IV procedure will then use just the true "signal" and recover a consistent estimate of the parameter (see Ashenfelter and Rouse 1998).[39]

We use data from Price (2012) on courts' no-look fee policies as a second measure of lawyer fees. The data were gathered via court websites or direct communication with the court. The data are not a good direct source for us because many fee policies changed between the period of our sample and when the data were collected in 2011. (Of the districts that reported when the no-look fee policy was set, 65% changed their policy between 2007 and 2011.) But it is a second measure of lawyer fees that is positively correlated with our fee data and so will serve as an excellent instrument correcting for measurement error.[40] Column 6 reports the results of this IV regression. The point estimate rises substantially to 0.46 and is statistically significant, suggesting measurement error may be attenuating our baseline estimate.[41] All told, we think the results of Table 3.5 suggest that lawyer fees do have a strong impact on chapter choice under a number of different plausible specifications of the independent variable.

In Table 3.6, we turn to a number of possible changes to our weighting, control variables, and sample selection. In our baseline specification we weight the data to make it representative of the bankrupt population. This is the most reasonable approach, but an

39. A second difficulty with our OLS specification is that although we maintain that Chapter 13 fees are largely exogenous, Chapter 7 fees are market-determined and therefore potentially endogenous. Our regression in column 5 accounts for this by simply excluding the Chapter 7 fee from the regression, but another approach is to instrument the price difference, p_d with a measure of the Chapter 13 court appointed fees. Column 6's IV specification thus will deal with this concern in addition to the measurement error problem discussed in the body of the text. As we discuss below, our instrument will not correct for omitted variables bias that may affect the Chapter 13 fee estimates.

40. Our fee data are highly correlated with the Price data but not identical, with a correlation of about 0.6.

41. The higher coefficient appears to largely be due to correcting for measurement error rather than endogenous Chapter 7 market prices, as we also considered regressions that used our own Chapter 13 fee data as an instrument for the fee gap and the coefficient was essentially the same as the baseline. It was only when we used either this secondary sample or, in unreported results, a split sample design that led to a substantial rise in the coefficient.

Table 3.6: *Robustness of State Fixed Effects Specification to Alternate Regression Specifications*

	Baseline	Weight by District	Alternate Covariate Specifications		Drop Out-of-state Filing Zips	Only Multi-district States	County Cultural Controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log 13/7 Fee Difference	0.32***	0.28***	0.33***	0.31***	0.32***	0.32***	0.37***
	[0.08]	[0.07]	[0.08]	[0.10]	[0.08]	[0.08]	[0.07]
Log Repayment Rate	0.33**	0.30**	0.29*		0.32*	0.32**	0.25*
	[0.16]	[0.14]	[0.15]		[0.17]	[0.16]	[0.15]
Log Monthly Income	-0.17	-0.18*			-0.19	-0.16	-0.14
	[0.12]	[0.10]			[0.13]	[0.13]	[0.13]
Log(Secured Debt/Total Debt)	-0.07	0.01	-0.21		-0.05	-0.08	-0.16
	[0.22]	[0.16]	[0.23]		[0.22]	[0.22]	[0.21]
Religion, Politics, and Crime Controls							yes
Observations	20,051	20,051	20,051	20,051	19,065	15,381	
R-squared	0.57	0.42	0.57	0.57	0.57	0.56	0.60

Notes: Each zip code in the regressions is weighted by the number of bankruptcies that occurred there in 2007, except column (2); see text for details. Demographic controls at the zip code level are included in each regression and are those summarized in Table 1: urban, population and population squared, marital status, household size, gender and age of head of household, education, race, age, employment status, self-employed, home ownership, and household income. State fixed effects are included in each regression. Religion and Politics Controls include the county's fraction adherents for six religious classifications (listed in the text), the fraction of the county that voted for George Bush in 2004, and the crime rate in 2004. Standard errors in brackets are clustered at the state level. * significant at 10%, ** significant at 5%, *** significant at 1%.

alternative would be to consider each district to be the relevant unit of analysis, as this is the level at which the variation occurs. As shown in Column 2, re-weighting the data so that each district is given equal weight gives essentially the same estimate.

Including a control for reported filer income forces the variation in repayment rates to come from different allowed expenses in the plan. But it may be that the relevant variation is better recovered by looking at the difference between reported income and expenses. Thus Column 3 drops the control for reported log income in the bankruptcy filing so that the repayment rate uses these changes in log income as a legitimate source of variation in repayment rates. Recall that we include extensive zip code–level controls for household income distributions, which we maintain in this regression. Using all the variation in repayment rates makes little difference to the price elasticity estimate, which is now 0.33.

Column 4 drops both repayment rates and reported income. This also has little effect on our estimate of the price elasticity, which is reassuring. The robustness of the point estimate suggests that pro-creditor policies captured by these controls are not heavily related to the allowed Chapter 13 fee. We return to this question later.

Although districts do not cross state boundaries by construction, a few zip codes along state borders have a large number of debtors crossing over state lines to file in a neighboring court that is more conveniently located. Thus these zip codes have non-matching states and districts and so, were they influential, could be affecting our results that use state fixed effects. Column 5 deals with this by dropping zip codes where most filings are outside the home state. We also dropped zip codes where filings were not predominantly in one district or another (less than 90%). This reduces our sample by about 5%, but has no substantive effect on the point estimate or standard error, suggesting that these zip codes are not affecting the results.

Since we use state fixed effects, our identification of the fee effect is dependent entirely on states with multiple districts. Single-district states provide information on zip-level coefficients but nothing else. Column 6 shows this to be the case by dropping all states with only one bankruptcy district. The fee elasticity is identical to the baseline.

Our last major concern is that Chapter 13 fees may be set by the court in a way that is correlated with other court policies encouraging Chapter 13 filing. Currently our demographic and state controls go a long way in controlling for differences across localities, but attitudes and other cultural factors at the local level may bias our results. For example, Chapter 13 may appeal to groups that feel it important to pay back one's obligations, and if these beliefs vary within a state but across districts, and are correlated with lawyer fees, this would confound our results.

We check for this by using county-level cultural data: the percent of a given county that voted for George Bush in the 2004 election; the percent of the population in seven religious categorizations (Mainline Protestant, Evangelical, Catholic, Orthodox, Other Christian, Jewish, and Islam); and the number of crimes reported in 2004 per 100,000 people (Association of Statisticians of American Religious Bodies, 1999–2001; U.S. Department of Commerce, 2000–2007). The regression results are reported in column 7. Note the sample size drops somewhat due to missing county-level data. Although many of these factors affected chapter choice, they did little to affect the price elasticity estimate, which rises to 0.37. Thus there is little evidence that unobserved cultural characteristics are biasing the results, and what evidence we do have suggest the bias is actually negative, attenuating our estimate.

Section 3.2.1 discussed the ambiguous relationship between a court's pro-creditor preferences and how it sets the no-look fee. Another angle on these preferences comes from Braucher (1999), who suggests that judges and trustees may approve higher fees to induce attorneys to write plans with high repayment rates for their clients, even if the filing household is unlikely to fulfill the requirements of such a demanding plan. If this occurs as part of district court policy, we should be able to observe a positive correlation between courts that approve higher fees and the repayment rates they require. If repayment rates are positive predictors of prices, this would raise concern that the price elasticity is also recovering unobserved district policies to stimulate Chapter 13 bankruptcies.

We look at this in Table 3.7, which performs a regression similar to our baseline speci-

Table 3.7: *Relationship Between District-set Chapter Price and District-set Repayment Rates*

	Log Chapter 13/7 Fee Difference (1)	Log Chapter 13 Price (2)
Log Repayment Rate	-0.43	-0.21
	[0.26]	[0.18]
Log Monthly Income	-0.24	0.05
	[0.26]	[0.10]
Log(Secured Debt/Total Debt)	0.13	0.16
	[0.28]	[0.25]
State Fixed Effects	yes	yes
R-squared	0.59	0.68

Notes: Each zip code in the regressions is weighted by the number of bankruptcies that occurred there in 2007. Demographic controls at the zip code level are included in each regression and are those summarized in Table 1: urban, population and population squared, marital status, household size, gender and age of head of household, education, race, age, employment status, self-employed, home ownership, and household income. N=25253. Standard errors in brackets are clustered at the state level.

fication, but with prices as the dependent variable. In column 1, the dependent variable is p_d, the log price gap between Chapter 7 and Chapter 13. In column 2, the dependent variable is the log Chapter 13 price. Both regressions include demographic controls and state fixed effects. In all specifications, the log repayment rate is *negatively*, rather than positively, related to the price variables and the coefficients are not significantly different from zero. Thus we find no evidence of district-level effects biasing our elasticity up. In fact, our results suggest that courts that want to help creditors (as proxied by their repayment rate policies) tend to mandate *lower* Chapter 13 lawyer fees. This is consistent with pro-creditor courts wanting to entice filers into Chapter 13 and transfer money from lawyers to non-priority creditors at the same time.

While this evidence suggests that high Chapter 13 fees are not proxying for pro-Chapter 13 court policies, unobserved correlation is a difficult question. It may be that some unaccounted for variation may be biasing our coefficient up. We think it unlikely that this bias is severe enough to overcome the substantial downward bias from measurement error suggested in Table 3.5.

3.4.3 Parameter Heterogeneity

The elasticity of 0.32 we consistently find may actually mask substantial heterogeneity across different locations or across different types of filers. For example, some state foreclosure regimes may make it easier for a lawyer to push a debtor between chapters by making an otherwise unpleasant Chapter 13 appear more pleasant to avoid foreclosure. Alternately, less educated clients may be more or less susceptible to the counsel of the lawyer. Of course, these things will also affect the optimal chapter choice for an agent (a highly educated worker may be more likely to have assets they wish to protect in Chapter 13) but what we are interested in is the extent to which the lawyer fee elasticity also varies with observable characteristics.

In Table 3.8, we look at how the fee elasticity varies based on a number of different characteristics at the zip code and state level. In each case, we simply re-estimate the baseline fixed effects specification while allowing the elasticity to be different for zip codes with different characteristics.[42] For example, in the first row we estimate separate elasticities for zips that are below and above the median high school graduation rate. Both estimates (0.33 and 0.31) are essentially identical to our baseline estimate, suggesting that high school graduation rates in an area are unrelated to the parameter of interest.

Similarly, median income and home ownership levels in a zip code do not interact at all with the fee elasticity estimate. Note that we are not saying that these variables have no effect on chapter choice. Rather we see little evidence that lawyer behavior differs across education, income, or home ownership groups. We do see some differences in the elasticity estimate across judicial foreclosure status and fraction of debt that is secured, which suggests that lawyers may respond more to price incentives when debtors' cases are more suited to Chapter 13. Unfortunately these estimates are quite imprecise. In the end, none of the

42. Methodologically, we split each characteristic into disjoint groups flagged by dummy variables, e.g., zips whose median income is below the national median, $LowInc_i$, vs above the national median, $HiInc_i$. Then we replace the elasticity term in Equation (3.2), δp_d, with $\delta_{Low}\left(p_d \times LowInc_i\right) + \delta_{Hi}\left(p_d \times HiInc_i\right)$.

Table 3.8: *Heterogeneity in Fee Effects*

	Below Median	At or Above Median	P-value on Test of Equality
High School Graduation Rate by Zip Code	0.33***	0.31***	0.7
	[0.09]	[0.08]	
Median Income by Zip Code	0.28***	0.34***	0.28
	[0.08]	[0.09]	
Homeownership by Zip Code	0.32***	0.32***	0.91
	[0.07]	[0.09]	
Judicial Foreclosure by State (0/1)	0.42***	0.18	0.15
	[0.11]	[0.11]	
Fraction Secured Debt by Bankruptcy District	0.23	0.44***	0.35
	[0.16]	[0.10]	

Notes: Dependent variable is Log Fraction 13. Each row repeats table 3, column 3 but interacts "Log 13/7 Fee Difference" with dummies for whether the zip code is above or below the sample median for the given variable. Each zip code in the regressions is weighted by 2007 population. Demographic controls at the zip code level are included in each regression and are those summarized in Table 1: urban, population and population squared, marital status, household size, gender and age of head of household, education, race, age, employment status, self-employed, home ownership, and household income. N=20,051. Standard errors in brackets are clustered at the state level. * significant at 10%, ** significant at 5%, *** significant at 1%.

differences in Table 3.8 are statistically significant, so we cannot reject the hypothesis that lawyers' response to their own price incentives is largely unaffected by their clients' needs or sophistication.

3.5 Conclusion

We estimate that a 10% rise in the Chapter 13 fee relative to Chapter 7 predicts a 3% switch from Chapter 7 to Chapter 13. This is true controlling for a wide array of socioeconomic factors, the pro-creditor legal precedents of a district court and unobserved differences across states. Although it is always possible that the estimated elasticity is biased by some correlation with an unobserved causal factor, we find little evidence to support this notion. With the exception of state policies, which we can purge with fixed effects, none of our controls for possible bias had a significant effect on our preferred point estimate. Furthermore, we find evidence suggesting that our baseline results are attenuated due to measurement error in the fee data, making the 3% response a lower bound. We also found that controlling for cultural and political differences strengthened, rather than weakened, our point estimate.

Chapter 13 fees vary enormously across district, from \$1500 to \$4400. Our estimates suggest that 5.4% of the cross-district differences in Chapter 13 rates likely could be eliminated by harmonizing Chapter 13 relative fees.[43] The statistical model as a whole does a good job of explaining cross-district variation in relative Chapter 13 rates; even without state fixed effects, the model can explain 73% of the variation across districts and 80% when we include the cultural controls in Table 3.6.

The results support the notion that lawyers can and do manipulate their client's filings to increase their revenues. Of course, we do not observe lawyer profits, so our conclusions would be tempered if it was found that in states with higher Chapter 13 prices, Chapter

43. To get this number, we estimate the change in district-level variation in the raw sample, then compare it with the variation that would exist if all fee gaps were were harmonized, using the baseline specification. This requires having a consistent estimate of the causal price elasticity. Note that this is different than the regression R-squared as we are interested in district variation rather than zip code variation.

13 filings also required more work. We know of no evidence to support this and Chapter 13 procedures appear to be fairly standardized, so we doubt this is a large problem given the enormous variation in Chapter 13 prices. Nevertheless, it could be worth examining in future research.

While we feel confident in our positive results, it is more difficult to make normative statements about the optimal Chapter 13 rate. Even if we were sure that being pushed into Chapter 13 is harmful to most debtors, this is certainly not true for all debtors.[44] Furthermore, the debtors are filing in order to avoid repaying their creditors, thus their gain may be the creditor's loss, which could result in higher equilibrium interest rates for other borrowers. Though we have no definitive answer, we consider a few of the larger welfare issues for debtors, creditors, and lawyers.

Since marginal filers use Chapter 13 when the Chapter 13 price is high and Chapter 7 when the Chapter 7 is relatively high, we know that they are paying relatively more even if they were indifferent to the value of the underlying service. If we believe lawyers are altering their signal in response to potential economic rents when the signal would otherwise be optimal for the debtor's case, this directly implies marginal filers are worse off as a result. Thus there is a cost for the marginal debtor to the lawyer's profit seeking behavior in some districts, if not all. In a competitive model with full information, where buyers directly bore the cost of their purchases, we would know that the optimal price level would be the unregulated market price. In this case it is not immediately obvious how one could empirically recover what the best Chapter 13 price is.[45]

44. Dobbie and Song find some suggestive evidence that Chapter 7 *dismissal* may not be as damaging as Chapter 13 *dismissal*. Of course, the differential damages from the two kinds of dismissal is probably not a good proxy for determining the relative value of successful completion of one chapter versus the other. Further, with so few Chapter 7 dismissals in their sample (only 2%) Dobbie and Song themselves shy away from assigning too much to the result.

45. Other than those discussed in the text, there are several potential benefits to the debtor from a Chapter 13 versus a Chapter 7. The temporary stay on debt collection is probably quite valuable to some debtors, though many debtors would probably be just as happy or happier removing the debts entirely under Chapter 7. For many debtors there may be substantial differences in the emotional costs of two filings. The drawn-out Chapter 13 process may be more draining than the relatively quick Chapter 7 filing, though the 13 may appeal to those with a preference for repaying at least some of their obligations.

There are also possible differential effects on future credit access. Most people filing bankruptcy are likely

Besides the fact that the debtors pay more as the price rises, a major concern for them is that a Chapter 13 plan is far more likely to be dismissed. If we had disaggregated zip code data on dismissals we could look at how changing fees changed dismissals in each district. One could imagine the effect going either way. If, for example, lawyers who file more Chapter 13s are better at filing them and so develop better cases for filers then this might, on the margin, overcome the much higher average dismissal rates of Chapter 13 filings vs. Chapter 7 filings.[46] Unfortunately, the only data we have available on dismissals is at the county level. Unreported results suggest that dismissal rates among Chapter 13 filings do not change with higher prices, though the estimates are sufficiently noisy that one cannot rule out some effect.

From the perspective of creditors, Chapter 13 filings introduce a new debt owed to the lawyer that is given priority over the unsecured creditor, thus reducing the odds of receiving payment if the plan fails. Of course, unsecured creditors receive almost nothing in a typical Chapter 7 filing and can only receive more under Chapter 13 if the debtor completes his plan, so it seems unlikely they are made worse off by moving debtors from 7 to 13. Thus moving debtors to Chapter 13 may be a net gain for them. Indeed, in equilibrium benefits or costs to creditors may substantially revert to future debtors in the form of higher or lower interest rates or constraints on borrowing.

One substantial difference between the two filings is the larger transfer of money from the debtor to the lawyer whenever the repayment plan fails. On average, this will be a regressive income redistribution, and so not something the court has a strong interest in supporting for its own sake.

We could mitigate the principal–agent problem by better aligning the incentives of

to already have fairly low credit scores, thus the marginal effect of actually filing may not be that substantial. Furthermore, filing a Chapter 7 assures that the debtor will be unable to file a further Chapter 7 for eight years and their debt burden has been eliminated. A Chapter 13 filing, on the other hand, signals to creditors that the filer may well be heading for a future Chapter 7 filing. Consistent with this, anecdotal evidence from speaking with bankruptcy attorneys suggests that Chapter 7 filings may actually be better for future credit than Chapter 13 filings.

46. Using this kind of model, Chandra and Staiger (2007) find evidence for benefits to patients who use services physicians have specialized in.

the lawyer and debtor. One mechanism that might increase the quality of information lawyers give would be for courts to further stretch disbursement to lawyers over the life of the repayment plan, thus putting them in a position more similar to the unsecured creditors. One expects that this will make them more concerned about whether the debtor will complete the plan. Alternatively, the required credit counseling component of post-2005 bankruptcies, or something like it, could perhaps serve as an alternative information source about chapter choice. These changes, though, might not be optimal due to the costs imposed on the creditors and, consequently, on the interest rates faced by future debtors.

Whatever the optimal balance of Chapter 13 and Chapter 7 bankruptcies, the larger issue remains. The evidence suggests that, as one might expect, in markets for goods sufficiently complex to require outside advice, that advice may well be compromised when coming from sellers with a financial stake in the decision.

Bibliography

Administrative Office of the United States Courts, 2012. "Report of Statistics Required by the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005." Washington, D.C.: 2013.

Aldy, Jospeh E., and W. Kip Viscusi. 2008. "Adjusting the Value of a Statistical Life for Age and Cohort Effects." *Review of Economics and Statistics* 90 (3): 573–581.

Anderson, Michael L., 2015. "As the Wind Blows: The Effects of Long-term Exposure to Air Pollution on Mortality." NBER Working Paper #21578.

Ashenfelter, Orley, and David Bloom. 1993. "Lawyers as Agents of the Devil in a Prisoner's Dilemma Game." NBER Working Paper #4447.

Ashenfelter, Orley, and Cecilia Rouse. 1998. "Income, Schooling, and Ability: Evidence from a New Sample of Identical Twins." *Quarterly Journal of Economics* 113 (1): 253–284.

Association of Statisticians of American Religious Bodies, 1999–2001. "Religious Congregations and Membership Study (computer file)." Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

Avery, Christopher, and Parag A. Pathak. 2015. "The Distributional Consequences of Public School Choice." NBER Working Paper #21525.

Banzhaf, H. Spencer, 2015. "Panel Data Hedonics: Rosen's First Stage and Difference-in-Differences as "Sufficient Statistics"." NBER Working Paper #21485.

Banzhaf, H. Spencer, and Randall P. Walsh. 2008. "Do People Vote with Their Feet? An Empirical Test of Tiebout's Mechanism." *American Economic Review* 98 (3): 843–863.

Bartik, Timothy J., 1987. "The Estimation of Demand Parameters in Hedonic Price Models." *Journal of Political Economy* 95 (1): 81–88.

Bayer, Patrick, Nathaniel Keohane, and Christopher Timmins. 2009. "Migration and hedonic valuation: The case of air quality." *Journal of Environmental Economics and Management* 58:1–14.

Bento, Antonio, 2013. "Equity Impacts of Environmental Policy." *Annual Review of Resource Economics* 5:181–196.

Bento, Antonio, Matthew Freedman, and Corey Lang. 2015. "Who Benefits from Environmental Regulation? Evidence from the Clean Air Act Amendments." *Review of Economics and Statistics* 97 (3): 610–622.

Black, Sandra E., 1999. "Do Better Schools Matter? Parental Valuation of Elementary Education." *Quarterly Journal of Economics* 114 (2): 577–599.

Borenstein, Seth, 2015. "AP analysis: Dozens of deaths likely from VW pollution dodge." October 3. http://bigstory.ap.org/article/a6925f0af82e44aaa1a1ed4b55d030f6, *Associated Press.*

Borenstein, Severin, 2002. "The Trouble with Electricity Markets: Understanding California's Restructuring Disaster." *Journal of Economic Perspectives* 16 (1): 191–211.

Braucher, Jean, 1993. "Lawyers and Consumer Bankruptcy: One Code, Many Cultures." *American Bankruptcy Law Journal* 67 (4): 501–583.

———. 1997. "Counseling Consumer Debtors to Make Their Own Informed Choices: A Question of Professional Responsibility." *American Bankruptcy Institute Law Review* 5:165–200.

———. 1999. "Options in Consumer Bankruptcy: An American Perspective." *Osgoode Hall Law Journal* 37:155–170.

———. 2001. "An Empirical Study of Debtor Education in Bankruptcy: Impact on Chapter 13 Completion Not Shown." *Am. Bankr. Inst. L. Rev.* 9:557.

Brunnermeier, Markus K., and Jonathan A. Parker. 2005. "Optimal Expectations." *American Economic Review* 112 (4): 1092–1118.

Busse, Meghan R., Devin G. Pope, Jaren C. Pope, and Jorge Silva-Risso. 2012. "Projection Bias in the Car and Housing Markets." NBER Working Paper #18212.

———. 2015. "The Psychological Effect of Weather on Car Purchases." *Quarterly Journal of Economics* 130 (1): 371–414.

Cellini, Stephanie Riegg, Fernando Ferreira, and Jesse Rothstein. 2010. "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design." *Quarterly Journal of Economics* 125 (1): 215–261.

Chandra, Amitabh, and Douglas O. Staiger. 2007. "Testing A Roy Model With Productivity Spillovers: Evidence From the Treatment of Heart Attacks." *Journal of Political Economy* 115 (1): 103–140.

Chay, Kenneth Y., and Michael Greenstone. 2005. "Does Air Quality Matter? Evidence from the Housing Market." *Journal of Political Economy* 113 (2): 376–424.

Chevalier, Judith, and Glenn D. Ellison. 1997. "Risk Taking by Mutual Funds as a Response to Incentives." *Journal of Political Economy* 114 (2): 389–432.

Chitano, P., J.J. Hosselet, C.E. Mapp, and L.M. Fabbri. 1995. "Effect of oxidant air pollutants on the respiratory system: insights from experimental animal research." *European Respiratory Journal* 8:1357–1371.

Cimorelli, Alan J., Steven G. Perry, Akula Venkatram, Jeffrey C. Weil, Robert J. Paine, Robert B. Wilson, Russell F. Lee, Warren D. Peters, and Roger W. Brode. 2005. "AERMOD: A dispersion model for industrial source applications. Part I: General model formulation and boundary layer characterization." *Journal of Applied Meteorology* 44 (5): 682–693.

Coglianese, John, Lucas W. Davis, Lutz Kilian, and James H. Stock. 2015. "Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand." NBER Working Paper #20980.

Conley, Timothy G., 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics* 92:1–45.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker. 2015. "Do Housing Prices Reflect Environmental Health Risks? Evidence from More than 1600 Toxic Plant Openings and Closings." *American Economic Review* 105 (2): 678–709.

Currie, Janet, Joshua Graff Zivin, Jamie Mullins, and Matthew Neidell. 2014. "What Do We Know About Short- and Long-Term Effects of Early-Life Exposure to Pollution?" *Annual Review of Resource Economics* 6 (1): 217–247.

Currie, Janet, and Matthew Neidell. 2005. "Air pollution and infant health: What can we learn from California's recent experience?" *Quarterly Journal of Economics* 120 (3): 1003–1030.

Currie, Janet, and Reed Walker. 2011. "Traffic Congestion and Infant Health: Evidence from EZPass." *American Economic Journal: Applied Economics* 3 (1): 65–90.

Davis, Lucas, 2004. "The Effect of Health Risk on Housing Values: Evidence from a Cancer Cluster." *American Economic Review* 94 (5): 1693–1704.

———. 2011. "The Effect of Power Plants on Local Housing Values and Rents." *The Review of Economics and Statistics* 93 (4): 1391–1402.

Dawsey, Amanda E., and Lawrence M. Ausubel. 2004. "Informal Bankruptcy." Working paper, Department of Economics, University of Maryland.

Dobbie, Will, and Jae Song. Forthcoming. "Debt Relief and Debtor Outcomes: Measuring the Effects of Consumer Bankruptcy Protection." *American Economic Review.*

Epple, Dennis, 1987. "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products." *Journal of Political Economy* 95 (1): 59–80.

Epple, Dennis, Radu Filimon, and Thomas Romer. 1984. "Equilibrium among Local Jurisdictions: Toward an Integrated Approach of Voting and Residential Choice." *Journal of Public Economics* 24:281–308.

———. 1993. "Existence of voting and housing equilibrium in a system of communities with property taxes." *Regional Science and Urban Economics* 23:585–610.

Epple, Dennis, and Holger Sieg. 1999. "Estimating Equilibrium Models of Local Jurisdictions." *Journal of Political Economy* 107 (4): 645–81.

Eraslan, Hulya, Wenli Li, and Pierre-Daniel Sarte. 2007. "The Anatomy of U.S. Personal Bankruptcy under Chapter 13." Working paper 07-31, Federal Reserve Bank of Philadelphia.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. "What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program." *American Economic Review* 102 (2): 965–993.

Fuchs, Victor R., 1978. "The Supply of Surgeons and the Demand for Operations." *Journal of Human Resources* 13:35–56.

Garrison, Mark, 2015. "What if it's not just Volkswagen?," *Marketplace.* http://www.marketplace.org/topics/sustainability/what-if-its-not-just-volkswagen. September 24.

Glaeser, Edward L., and Joshua D. Gottlieb. 2008. "The Economics of Place-Making Policies." *Brookings Papers on Economic Activity* 1:155–253.

Graff Zivin, Joshua, and Matthew Neidell. 2012. "The Impact of Pollution on Worker Productivity." *American Economic Review* 102 (7): 3652–3673.

Grainger, Corbett A., 2012. "The distributional effects of pollution regulations: Do renters fully pay for cleaner air?" *Journal of Public Economics* 96:840–852.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2012. "Liqbidity Constraints and Consumer Bankruptcy: Evidence from Tax Rebates." *Review of Economics and Statistics* 96 (3): 431–443.

Gruber, Jonathan, and Maria Owings. 1996. "Physician Financial Incentives and Cesarean Section Delivery." *Rand Journal of Economics* 27 (1): 99–123.

Halla, Martin, 2007. "Divorce and the Excess Burden of Lawyers." IZA Discussion Papers #2962.

Hanna, Rema, and Paulina Oliva. 2015. "The Effect of Pollution on Labor Supply: Evidence from a Natural Experiment in Mexico City." *Journal of Public Economics* 122:68–79.

Harrington, David E., and Kathy J. Krynski. 2002. “The Effect of State Regulations on Cremation Rates: Testing for Demand Inducement in Funeral Markets.” *Journal of Law and Economics* 45 (1): 199–226.

Hu, Shishan, Scott Fruin, Kathleen Kozawa, Steve Mara, Suzanne E. Paulson, and Arthur M. Winer. 2009. “A wide area of air of air pollutant impact downwind of a freeway during pre-sunrise hours.” *Atmospheric Environment* 43 (16): 2541–49.

Karlsson, Martin, Maike Schmitt, and Nicolas R. Ziebarth. 2015. “The Short-term Population Health Effects of Weather and Pollution.” Working Paper.

Karner, Alex A., Douglas S. Eisinger, and Deb A. Niemeier. 2010. “Near-Roadway Air Quality: Synthesizing the Findings from Real-World Data.” *Environmental Science & Technology* 44 (14): 5334–5344.

Kelejian, Harry H., and Ingmar R. Prucha. 2007. “HAC estimation in a spatial framework.” *Journal of Econometrics* 140:131–154.

Kline, Patrick, 2010. “Place Based Policies, Heterogeneity, and Agglomeration.” *American Economic Review* 101 (3): 532–537.

Knittel, Christopher R., Douglas L. Miller, and Nicholas J. Sanders. 2014. “Caution, drivers! Children present: Traffic, pollution, and infant health.” Working Paper.

Kuminoff, Nicolai V., and Jaren C. Pope. 2014. “Do “Capitalization Effects” for Public Goods Reveal the Public’s Willingness to Pay?” *International Economic Review* 55 (4): 1227–50.

Kuminoff, Nicolai V., V. Kerry Smith, and Christopher Timmins. 2013. “The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets.” *Journal of Economic Literature* 51 (4): 1007–1062.

Layton, Timothy, Frank McIntyre, and Daniel M. Sullivan. 2010. “Did BAPCPA Deter the Wealthy? The 2005 Bankruptcy Reform’s Effect on Filings Across the Income and Asset Distribution.” Working Paper.

Lefgren, Lars, and Frank McIntyre. 2009. “Explaining the Puzzle of Cross-State Differences in Bankruptcy Rates.” *Journal of Law and Economics* 52 (2): 367–393.

Lefgren, Lars, Frank McIntyre, and Michelle Miller. 2010. “Chapter 7 or 13: Are Client or Lawyer Interests Paramount?” *The B.E. Journal of Economic Analysis and Policy* 10 (1): 1–44.

Lehnert, Andreas, and Dean M. Maki. 2002. “Consumption, Debt and Portfolio Choice: Testing the Effect of Bankruptcy Law.” Federal Reserve Bank Working Paper #02-14.

Levitt, Steven D., and Chad Syverson. 2008. “Market Distortions when Agents Are Better Informed: The Value of Information in Real Estate Transactions.” *The Review of Economics and Statistics* 90 (4): 599–611.

Li, Wenli, Michelle J. White, and Ning Zhu. 2011. “Did Bankruptcy Reform Cause Mortgage Defaults to Rise?” *American Economic Journal: Economic Policy* 3 (4): 123–47.

Linden, Leigh, and Jonah E. Rockoff. 2008. “Estimates of the Impact of Crime Risk on Property Values from Megan’s Laws.” *American Economic Review* 98 (3): 1103–1127.

Lleras-Muney, Adriana, 2010. “The Needs of the Army: Using Compulsory Relocation of the Military to Estmate the Effect of Air Pollutants on Children’s Health.” *Journal of Human Resources* 45 (3): 549–590.

Luechinger, Simon, 2014. “Air pollution and infant mortality: A natural experiment from power plan desulfurization.” *Journal of Health Economics* 37:219–231.

Mann, Ronald J., and Katherine M. Porter. 2010. “Saving Up for Bankruptcy.” *Georgetown Law Journal* 98:289–339.

Mayraz, Guy, 2013. “Wishful Thinking.” Working Paper.

Mian, Atif R., Amir Sufi, and Francesco Trebbi. 2014. “Foreclosures, House Prices, and the Real Economy.” Chicago Booth Research Paper No. 13-41.

Mian, Atif, and Amir Sufi. 2009. “The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis.” *Quarterly Journal of Economics* 124 (4): 1449–1496.

Miguel, Edward, and Michael Kremer. 2004. “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities.” *Econometrica* 72 (1): 159–217.

Montiel Olea, Jose Luis, and Carolin Pflueger. 2013. “A Robust Test for Weak Instruments.” *Journal of Business and Economic Statistics* 31 (3): 358–369.

Moretti, Enrico, and Matthew Neidell. 2011. “Pollution, Health, and Avoidance Behavior: Evidence from the Ports of Los Angeles.” *Journal of Human Resources* 46 (1): 154–75.

Neidell, Matthew, 2009. “Information, Avoidance Behavior, and Health: The Effect of Ozone on Asthma Hospitalizations.” *Journal of Human Resources* 44 (2): 450–478.

Neidell, Matthew J., 2004. “Air pollution, health, and socio-economic status: the effect of outdoor air quality on childhood asthma.” *Journal of Health Economics* 23 (6): 1209–1236.

Neustadter, Gary, 1986. “When Lawyer and Client Meet: Observations of Interviewing and Counseling Behavior in the Consumer Bankruptcy Law Office.” *Buffalo Law Review* 35 (1): 177–284.

Norberg, Scott F., and Nadja Schreiber Compo. 2007. "Report on an empirical study of district variations, and the roles of judges, trustees and debtors' attorneys in Chapter 13 bankruptcy cases." *American Bankruptcy Law Journal* 81:101–158.

Norberg, Scott F., and Andrew J. Velkey. 2006. "Debtor Discharge and Creditor Repayment in Chapter 13." *Creighton Law Review* 39:473–560.

Palmquist, Raymond B., 2005. "Property Value Models." Chap. 16 in *Handbook of Environmental Economics,* edited by Karl-Göran Mäler and Jeffery R. Vincent, 2:763–819. Amsterdam: North Holland.

Perry, Steven G., Alan J. Cimorelli, Robert J. Paine, Roger W. Brode, Jeffrey C. Weil, Akula Venkatram, Robert B. Wilson, Russell F. Lee, and Warren D. Peters. 2005. "AERMOD: A dispersion model for industrial source applications. Part II: Model performance against 17 field study databases." *Journal of Applied Meteorology* 44 (5): 694–708.

Pope, C. Arden, Richard T. Burnett, Michael J. Thun, Eugenia E. Calle, Daniel Krewski, Kazuhiko Ito, and George D. Thurston. 2002. "Lung Cancer, Cardiopulmonary Mortality, and Long-term Exposure to Fine Particulate air Pollution." *Journal of the American Medical Association* 287 (6): 1132–1141.

Pope, Devin G., Jaren C. Pope, and Justin R. Sydnor. 2014. "Focal Points and Bargaining in Housing Markets." Working paper.

Pope, Jaren C., 2008. "Fear of crime and housing prices: Household reactions to sex offender registries." *Journal of Urban Economics* 64:601–614.

Pouliot, Guillame, 2015. "Spatial Econometrics for Misaligned Data." Working Paper.

Price, Bruce M., 2012. "No look attorneys' fees and the attorneys who are looking: An empirical analysis of presumptively approved attorneys' fees in ch. 13 bankruptcies and a proposal for reform." *American Bankruptcy Institute Law Review* 20:291–351.

Ridker, Ronald G., and John A. Henning. 1967. "The Determinants of Residential Property Values with Special Reference to Air Pollution." *Review of Economics and Statistics* 49 (2): 256–257.

Rosen, Sherwin, 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1): 34–55.

Schlenker, Wolfram, and W. Reed Walker. Forthcoming. "Airports, Air Pollution, and Contemporaneous Health." *Review of Economic Studies.*

Sieg, Holger, V. Kerry Smith, H. Spencer Banzhaf, and Randy Walsh. 2004. "Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods." *International Economic Review* 45 (4): 1047–1077.

Sillman, Sanford, 1999. "The relation between ozone, NO_x, and hydrocarbons in urban and polluted rural environments." *Atmospheric Environment* 33:1821–1845.

Smith, V. Kerry, and Ju-Chin Huang. 1995. "Can Markets Value Air Quality? A Meta-Analysis of Hedonic Property Value Models." *Journal of Political Economy* 103 (1): 209–227.

South Coast Air Quality Management District, 2000. *Review of RECLAIM Findings.* Internal Report, October 20.

———. 2015a. Public Records Requests #80085 and #80086.

———. 2015b. Correspondence in reply to Public Records Request #80089.

Stock, James H., Jonathan H. Wright, and Motohiro Yogo. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business and Economic Statistics* 20 (4): 518–529.

Stock, James H., and Motohiro Yogo. 2002. "Testing for Weak Instruments in Linear IV Regression." NBER Technical Working Paper #284.

Sullivan, Ralph J., 1969. *Preliminary air pollution survey of odorous compounds, A literature review.* U.S. National Air Pollution Control Administration Publication No. APTD 69-42.

Sullivan, Teresa A., Elizabeth Warren, and Jay Lawrence Westbrook. 1988. "Laws, Models, and Real People: Choice of Chapter in Personal Bankruptcy." *Law & Social Inquiry* 13 (4): 661–706.

———. 1994. "The Persistence of Local Legal Culture: Twenty Years of Experience From the Federal Bankruptcy Courts." *Harvard Journal of Law and Public Policy* 17 (3): 801–865.

Tiebout, Charles M., 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 64 (5): 416–424.

U.S. Courts, 2013. "Bankruptcy Statistics." U.S. Federal Courts Website. http://www.uscourts.gov/Statistics/BankruptcyStatistics.aspx.

U.S. Department of Commerce, Bureau of the Census, Population Estimates Program, 2000–2007. "County Population Estimates and Estimated Components of Change (computer file)." Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

U.S. National Archives and Records Administration, 2004. *Code of Federal Regulations.* Title 40. Protection of Environment.

Van den Steen, Eric, 2004. "Rational Overoptimism (and Other Biases)." *American Economic Review* 94 (4): 1141–1151.

Walker, W. Reed, 2013. "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce." *Quarterly Journal of Economics* 128 (4): 1787–1835.

Weare, Christopher, 2003. *The California Electricity Crisis: Causes and Policy Options.* San Fransisco, CA: Public Policy Institute of California.

White, Michelle J., and Ning Zhu. 2010. "Saving Your Home in Chapter 13 Bankruptcy." *The Journal of Legal Studies* 39 (1): 33–61.

Woodruff, Tracey J., Jennifer D. Parker, and Kenneth C. Schoendorf. 2006. "Fine Particulate Matter ($PM_{2.5}$) Air Pollution and Selected Causes of Postneonatal Infant Mortality in California." *Environmental Health Perspectives* 114 (5): 786–790.

Zhu, Ning, 2011. "Household Consumption and Personal Bankruptcy." *The Journal of Legal Studies* 40 (1): 1–37.

Appendix A

Appendix to Chapter 1

A.1 Firm Data Construction

A.1.1 Geocoding

The accurate geocoding of pollution sources is obviously critical when analyzing the effect these sources have on the surrounding population. Administrative records on the latitude and longitude of each smoke stack operated by the firm would be the ideal data. Regulators often collect this data for the explicit purpose of dispersion modeling, and though SCAQMD does collect this data, they are unavailable for public use (SCAQMD 2015b). In lieu of direct geographic data for each smoke stack, I follow the literature and simply geocode the firms' street addresses, taking care to use the actual operating address of the firm and not a corporate or mailing address which are often listed in databases. For large firms and firms that match to interpolated street addresses instead of parcel centroids, I double-checked the coordinates using satellite photos from Google Maps to make sure the geographic point that represents the firm is reasonably close to the actual smoke stacks.[1]

1. This is potentially important because the firm's "store-front" address right on the street is often at the edge of the property, far away from the smoke stacks. Using unchecked street addresses can introduce significant errors (1–2 km) for firms that occupy large parcels of land.

A.1.2 Facility ID

SCAQMD assigns each facility an ID number; however, a facility may have more than one ID number in the data, both over time and cross-sectionally. This is primarily a concern when matching firms to the NEI, as described in Section A.1.3. It might also affect the pattern of firm behavior described by Figure 1.5, though this figure is only descriptive and not used in any calculations.

A facility's ID can change under a number of circumstances: the facility is sold, changes its name, or some part of its address changes. For the most part, these changes occur for superficial reasons, e.g., a zip code or street suffix is changed. To construct unique facility ID's, I flagged every pair of facilities less than 400 meters apart and visually inspected satellite photos and emissions data for every cluster of neighboring facilities. First, firms were merged if they occupied the same or neighboring parcels and shared breaks in their time series of emissions. For example, Facility A emits 25 tons per quarter from 1994 to 1999Q3 and then is missing from the data, while Facility B, located at the same parcel of land as A, enters the data in 1999Q4 and begins emitting 25 tons per quarter. Facilities were also merged if they had similar names and occupied the same or neighboring parcels of land. These merges were verified by checking whether or not the firms appeared separately in the NEI.

A.1.3 Stack Data from the NEI

Data for each firm's smoke stacks is taken from the National Emissions Inventory (NEI) from 1999 and 2002. Besides the smoke stack parameters, the NEI also has data on firm's name, address, SIC, and the equipment's SSC, and the estimated emissions by pollutant for each stack.[2] It also includes the ID number assigned to the facility by state-level regulators. For SCAQMD firms, this "state ID" consists of a county code, an air basin code, an air

2. The Source Classification Codes (SCC) for point pollution sources are a hierarchical index used by the EPA that categorize pollution-generating equipment by combustion type, fuel type, and size. It is analogous to the hierarchical SIC and NAICS industry codes.

district code, and the SCAQMD-assigned facility ID. Using this reconstructed ID, I was able to match most facilities in the SCAQMD emissions data to the NEI using either their own facility ID or an ID from a facility I had previously matched to it as described in section A.1.2. I used the 2002 NEI data whenever possible, falling back to the 1999 database when necessary. For facilities whose ID's did not match either dataset, I tried to match them using firm address and name. Firms that still did not match were almost all small firms that had ceased to exist before the NEI 1999 data was collected. These firms should have little impact on the overall results and were dropped. For matched facilities, I verified that individual stacks were not duplicated.

Many of the stack parameters in the NEI are flagged as imputed values. The imputation process was not well documented, so I re-imputed them using the median stack parameters from all non-imputed stacks in the SIC and SCC group. Finally, when passing the stack parameters to AERMOD, I weighted each stack according to its reported emissions in the NEI.

A.2 Supplementary Figures and Tables

Table A.1: *Pollution and House Prices, Full Diff-in-diff Results*

	(1) ln Price	(2) Aermod	(3) ln Price
Aermod_pre×post	0.0032***	-0.4328***	
	[0.0008]	[0.0748]	
Aermod			-0.0073***
			[0.0024]
Year-Quarter Effects			
1997Q2	0.0712	-0.3588	0.0686
	[0.0681]	[0.4045]	[0.0679]
1997Q3	0.1384	0.4692	0.1419
	[0.1310]	[0.7743]	[0.1309]
1997Q4	0.2033	-0.7367	0.1979
	[0.1885]	[1.1304]	[0.1882]
1998Q1	0.2578	-0.5633	0.2536

Table A.1: *(continued)*

	[0.2447]	[1.4530]	[0.2444]
1998Q2	0.3393	-2.9441*	0.3177
	[0.2978]	[1.7688]	[0.2973]
1998Q3	0.4252	-0.6891	0.4201
	[0.3466]	[2.0542]	[0.3462]
1998Q4	0.4730	-0.6668	0.4681
	[0.3920]	[2.3262]	[0.3916]
1999Q1	0.5225	-1.9205	0.5084
	[0.4344]	[2.5840]	[0.4339]
1999Q2	0.5827	-2.0290	0.5678
	[0.4728]	[2.8160]	[0.4722]
1999Q3	0.6537	-0.8812	0.6473
	[0.5086]	[3.0291]	[0.5082]
1999Q4	0.6790	-1.1694	0.6704
	[0.5420]	[3.2295]	[0.5416]
2000Q1	0.7352	-2.1188	0.7197
	[0.5700]	[3.4038]	[0.5694]
2000Q2	0.7930	-2.2088	0.7767
	[0.5957]	[3.5583]	[0.5951]
2000Q3	0.8422	-1.7783	0.8291
	[0.6178]	[3.7013]	[0.6172]
2000Q4	0.8985	-2.7325	0.8784
	[0.6366]	[3.8267]	[0.6360]
2001Q1	0.9363	-0.2790	0.9342
	[0.6529]	[3.9004]	[0.6526]
2001Q2	0.9765	-1.0513	0.9687
	[0.6657]	[3.9849]	[0.6654]
2001Q3	1.0221	-0.9022	1.0154
	[0.6758]	[4.0596]	[0.6754]
2001Q4	1.0844	-1.3753	1.0743
	[0.6822]	[4.1135]	[0.6819]
2002Q1	1.1269	-1.0652	1.1191
	[0.6865]	[4.1590]	[0.6862]
2002Q2	1.2030*	-1.3092	1.1934*
	[0.6878]	[4.1862]	[0.6875]
2002Q3	1.2714*	-0.9633	1.2643*
	[0.6870]	[4.2010]	[0.6868]
2002Q4	1.3142*	-0.7296	1.3088*
	[0.6829]	[4.2021]	[0.6828]
2003Q1	1.3676**	-0.5854	1.3633**
	[0.6768]	[4.1923]	[0.6768]
2003Q2	1.4447**	-0.6273	1.4401**
	[0.6689]	[4.1759]	[0.6689]

Table A.1: *(continued)*

2003Q3	1.4980**	0.0445	1.4983**
	[0.6587]	[4.1541]	[0.6589]
2003Q4	1.5687**	-0.0061	1.5687**
	[0.6479]	[4.1249]	[0.6481]
2004Q1	1.6395***	0.3494	1.6421***
	[0.6351]	[4.0926]	[0.6355]
2004Q2	1.7394***	0.5440	1.7434***
	[0.6225]	[4.0642]	[0.6230]
2004Q3	1.7928***	1.2637	1.8021***
	[0.6095]	[4.0432]	[0.6101]
2004Q4	1.8183***	1.4946	1.8293***
	[0.5975]	[4.0220]	[0.5982]
2005Q1	1.8874***	2.0931	1.9028***
	[0.5876]	[4.0156]	[0.5884]
2005Q2	1.9504***	2.1437	1.9661***
	[0.5793]	[4.0313]	[0.5801]
2005Q3	1.9984***	2.8258	2.0191***
	[0.5752]	[4.0567]	[0.5761]
2005Q4	2.0246***	3.4054	2.0496***
	[0.5762]	[4.1181]	[0.5771]
Demographic Time Trends			
Loan-to-Value Ratio $\times t$	0.0312	0.2497	0.0330
	[0.0305]	[0.1623]	[0.0304]
Loan-to-Value Ratio $\times t^2$	0.0003	-0.0222	0.0001
	[0.0032]	[0.0170]	[0.0032]
Interest Rate $\times t$	-0.0368	0.2035	-0.0353
	[0.0311]	[0.1916]	[0.0310]
Interest Rate $\times t^2$	0.0058*	-0.0212	0.0057*
	[0.0033]	[0.0202]	[0.0033]
log Median Income $\times t$	0.0068	-0.1038**	0.0061
	[0.0073]	[0.0412]	[0.0073]
log Median Income $\times t^2$	-0.0034***	0.0044	-0.0033***
	[0.0008]	[0.0039]	[0.0008]
Geographic Time Trends			
Grid 1 $\times t$	-0.0042	0.5283***	-0.0003
	[0.0175]	[0.0682]	[0.0175]
Grid 1 $\times t^2$	-0.0001	-0.0390***	-0.0004
	[0.0017]	[0.0052]	[0.0017]
Grid 2 $\times t$	0.0051	0.7910***	0.0109
	[0.0173]	[0.0821]	[0.0172]
Grid 2 $\times t^2$	-0.0018	-0.0657***	-0.0023
	[0.0017]	[0.0061]	[0.0017]
Grid 3 $\times t$	0.0102	1.0460***	0.0179

Table A.1: *(continued)*

	[0.0187]	[0.1474]	[0.0186]
Grid 3×t^2	-0.0029	-0.0806***	-0.0035*
	[0.0019]	[0.0129]	[0.0019]
Grid 4×t	0.0042	0.9010***	0.0109
	[0.0209]	[0.1276]	[0.0209]
Grid 4×t^2	-0.0012	-0.0800***	-0.0018
	[0.0022]	[0.0144]	[0.0022]
Grid 5×t	0.0592***	0.4166***	0.0623***
	[0.0185]	[0.0508]	[0.0185]
Grid 5×t^2	-0.0054***	-0.0307***	-0.0056***
	[0.0018]	[0.0046]	[0.0018]
Grid 6×t	-0.0861***	0.5024***	-0.0824***
	[0.0213]	[0.0488]	[0.0213]
Grid 6×t^2	0.0074***	-0.0364***	0.0071***
	[0.0021]	[0.0045]	[0.0021]
Grid 7×t	-0.0177	0.4356***	-0.0145
	[0.0188]	[0.1151]	[0.0188]
Grid 7×t^2	0.0015	-0.0237***	0.0013
	[0.0019]	[0.0074]	[0.0019]
Grid 8×t	-0.0459**	0.5502***	-0.0418**
	[0.0199]	[0.0675]	[0.0199]
Grid 8×t^2	0.0042**	-0.0311***	0.0040**
	[0.0020]	[0.0060]	[0.0020]
Grid 9×t	-0.0008	0.4396***	0.0024
	[0.0168]	[0.0574]	[0.0168]
Grid 9×t^2	0.0000	-0.0272***	-0.0002
	[0.0016]	[0.0049]	[0.0016]
Grid 10×t	-0.0000	0.4612***	0.0033
	[0.0175]	[0.0526]	[0.0175]
Grid 10×t^2	-0.0004	-0.0341***	-0.0006
	[0.0017]	[0.0050]	[0.0017]
Grid 11×t	0.0111	0.7269**	0.0164
	[0.0236]	[0.3109]	[0.0232]
Grid 11×t^2	-0.0015	-0.0266	-0.0017
	[0.0024]	[0.0196]	[0.0023]
Grid 12×t	0.0301*	0.7649***	0.0357**
	[0.0176]	[0.1242]	[0.0176]
Grid 12×t^2	-0.0027	-0.0540***	-0.0031*
	[0.0017]	[0.0076]	[0.0017]
Grid 13×t	-0.0132	0.7943***	-0.0074
	[0.0180]	[0.0958]	[0.0180]
Grid 13×t^2	-0.0002	-0.0476***	-0.0005
	[0.0018]	[0.0076]	[0.0018]

Table A.1: *(continued)*

Grid 14×t	0.1239*	1.3842**	0.1341*
	[0.0730]	[0.6399]	[0.0753]
Grid 14×t^2	-0.0141*	-0.1041	-0.0148*
	[0.0084]	[0.0637]	[0.0086]
Grid 15×t	-0.0038	0.7585***	0.0018
	[0.0205]	[0.0755]	[0.0205]
Grid 15×t^2	-0.0016	-0.0514***	-0.0020
	[0.0021]	[0.0075]	[0.0021]
Grid 16×t	0.0261	0.6306***	0.0307*
	[0.0176]	[0.1031]	[0.0176]
Grid 16×t^2	-0.0024	-0.0758***	-0.0030*
	[0.0017]	[0.0135]	[0.0017]
Grid 17×t	0.0127	0.4856***	0.0163
	[0.0183]	[0.0738]	[0.0183]
Grid 17×t^2	-0.0017	-0.0395***	-0.0020
	[0.0018]	[0.0074]	[0.0018]
Method	OLS	OLS	2SLS

Note: N=41,771. Table presents the full regression output Columns 2–4 of Table 1.5. Controls also include property fixed effects.

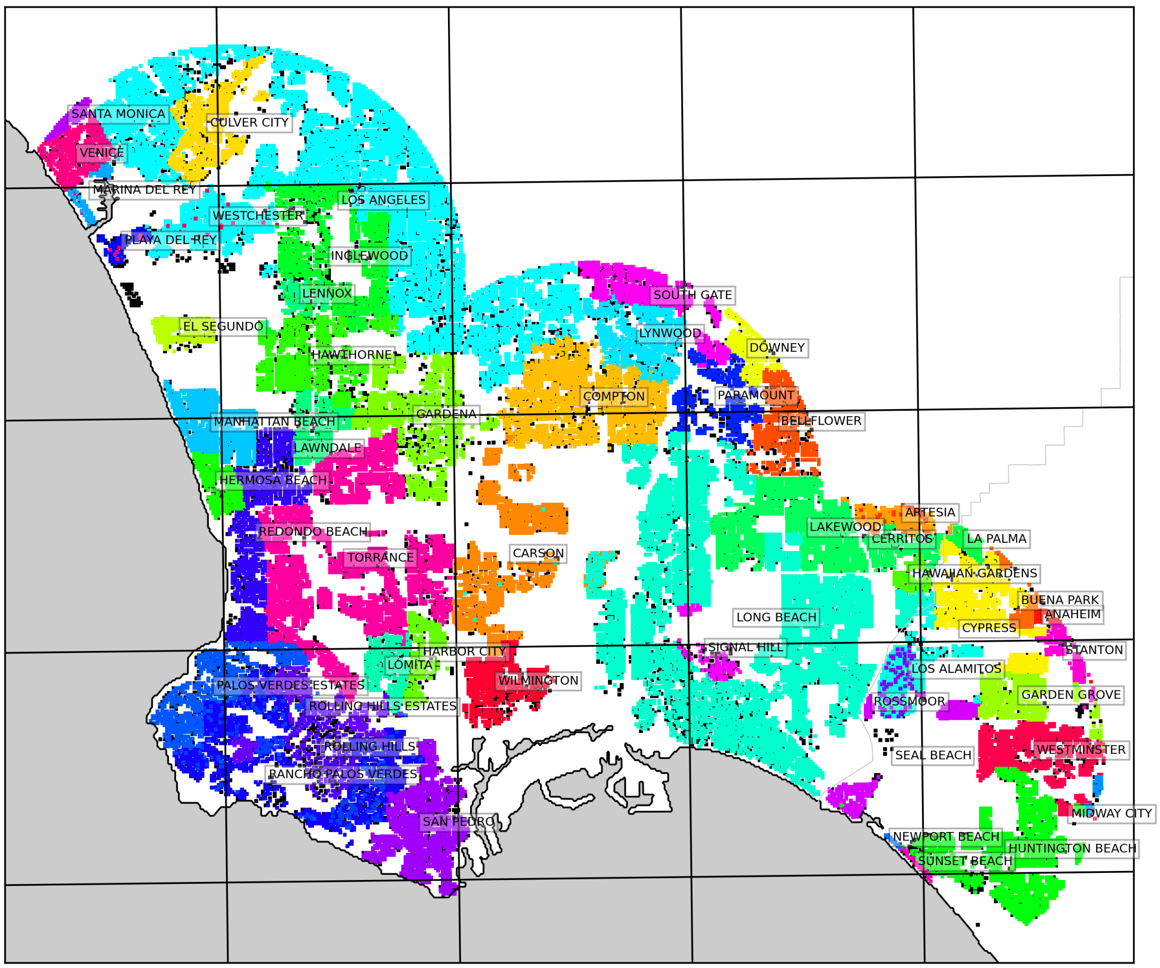

Figure A.1: *Cities in Sample Area with 10-km Grid*

Notes: Colors denote parcels belonging to different cities. Black parcels have no city data.

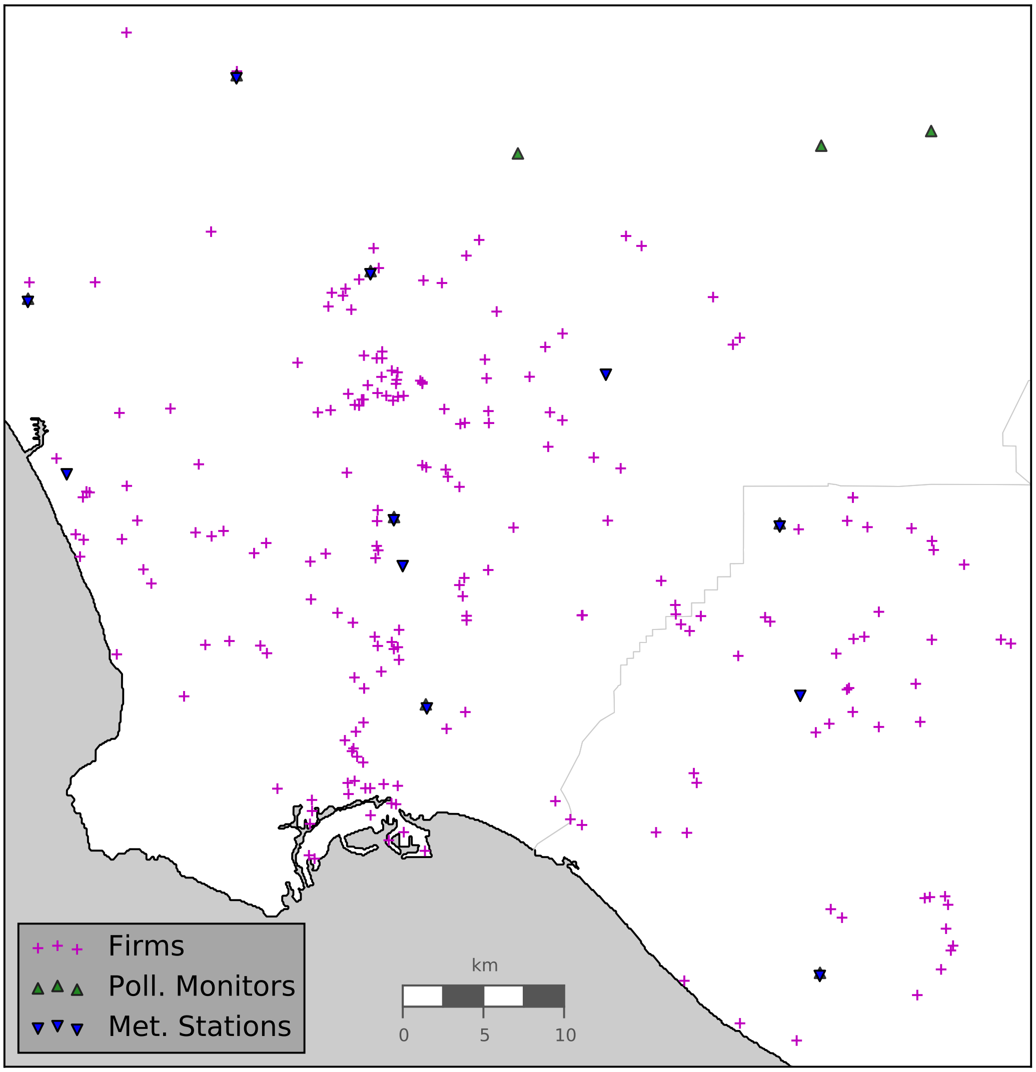

Figure A.2: *Monitoring Station and Firm Locations*

Notes: Firms and meteorology stations are restricted to those that contribute to the main regression sample. Pollution monitors restricted to those with constant NO_x coverage over 1997–2005.

Appendix B

Appendix to Chapter 2

B.1 Proofs

Definition 4 (Notation)**.** For simplicity of notation, let

$$X^j = \int_{\tilde{y}_{j-1}}^{\tilde{y}_j} h_p(p_j, y) f(y)\, \mathrm{d}y - S_p^j(p_j)$$

$$H^{ij} = h(p_i, \tilde{y}_j) f(\tilde{y}_j)$$

$$V^{ij} = V(\tilde{y}_i, p_j, g_j)$$

Lemma 1. *The following hold:*

$$H^{ij} > 0 \quad \forall i, j$$

$$X_j < 0 \quad \forall j$$

$$V_y^{j,j} - V_y^{j,j+1} < 0 \quad \forall j < J$$

Proof. $H^{ij} > 0$ follows from non-negative demand and non-negative probability distribution f. $X^j < 0$ follows from downward-slopping demand and upward-slopping supply: $h_p < 0$ and $S_p^j > 0$. For the last inequality, recall that that, by definition of $\tilde{y}_1$, we have $V(\tilde{y}_1, p_1, g_1) = V(\tilde{y}_1, p_2, g_2)$ and for $y \in (\tilde{y}_1, \tilde{y}_2]$, $j = 2$ is preferred to $j = 1$. Thus, for $0 < \varepsilon < \tilde{y}_2 - \tilde{y}_1$, we have $V(\tilde{y}_1 + \varepsilon, p_1, g_1) < V(\tilde{y}_1 + \varepsilon, p_2, g_2)$, so $V_y^{j,j} - V_y^{j,j+1} < 0$. $\square$

Lemma 2. *If V satisfies single-crossing, then, for arbitrary p, g, and $y_2 > y_1 > 0$,*

$$\frac{V_p(y_1, p, g)}{V_p(y_2, p, g)} > \frac{V_g(y_1, p, g)}{V_g(y_2, p, g)}$$

Proof. Single-crossing requires $M = -V_g/V_p$ to be monotonically increasing in y.

$$\begin{aligned} & M(y_2, p, g) > M(y_1, p, g) \\ \Leftrightarrow & -\frac{V_g(y_2, p, g)}{V_p(y_2, p, g)} > -\frac{V_g(y_1, p, g)}{V_p(y_1, p, g)} \\ \Leftrightarrow & \frac{V_p(y_1, p, g)}{V_p(y_2, p, g)} > \frac{V_g(y_1, p, g)}{V_g(y_2, p, g)} \end{aligned}$$

□

Proposition 1. *The following conditions hold:*

$$\frac{\partial \tilde{y}_2}{\partial g_2} > 0; \quad \frac{\partial p_2}{\partial g_2} > 0; \quad \frac{\partial p_3}{\partial g_2} < 0 \tag{2.3}$$

$$\frac{\partial \tilde{y}_1}{\partial g_2} \propto \frac{\partial p_1}{\partial g_2} \tag{2.4}$$

and the sign of $\partial \tilde{y}_1 / \partial g_2$ is ambiguous.

Proof. For $J = 3$, the equilibrium conditions given by Definition 3 can be written $\boldsymbol{F}(\boldsymbol{\theta}, \boldsymbol{g}) = \boldsymbol{0}$ where $\boldsymbol{\theta} = (\tilde{y}_1, \tilde{y}_2, p_1, p_2, p_3)$, $\boldsymbol{g} = (g_1, g_2, g_3)$, and

$$\boldsymbol{F}(\boldsymbol{\theta}, \boldsymbol{g}) = \begin{pmatrix} V(\tilde{y}_1, p_1, g_1) - V(\tilde{y}_1, p_2, g_2) \\ V(\tilde{y}_2, p_2, g_2) - V(\tilde{y}_2, p_3, g_3) \\ \int_{\underline{y}}^{\tilde{y}_1} h(p_1, y) f(y) \, \mathrm{d}y - S^1(p_1) \\ \int_{\tilde{y}_1}^{\tilde{y}_2} h(p_2, y) f(y) \, \mathrm{d}y - S^2(p_2) \\ \int_{\tilde{y}_2}^{\overline{y}} h(p_3, y) f(y) \, \mathrm{d}y - S^3(p_3) \end{pmatrix}$$

The Jacobian of $\boldsymbol{F}$ with respect to $\boldsymbol{\theta}$ is

$$\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\theta}} = \begin{pmatrix} V_y^{11} - V_y^{12} & 0 & V_p^{11} & -V_p^{12} & 0 \\ 0 & V_y^{22} - V_y^{23} & 0 & V_p^{22} & -V_p^{23} \\ H^{11} & 0 & X^1 & 0 & 0 \\ -H^{21} & H^{22} & 0 & X^2 & 0 \\ 0 & -H^{32} & 0 & 0 & X^3 \end{pmatrix}$$

By Lemma 1, the Jacobian determinant is strictly negative.

$$\begin{aligned} \det \frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\theta}} = {} & H^{21}H^{32}V_p^{12}V_p^{23}X^1 + H^{11}H^{32}V_p^{11}V_p^{23}X^2 + H^{11}H^{22}V_p^{11}V_p^{22}X^3 \\ & + \left(X^1X^2X^3(V_y^{22} - V_y^{23}) - H^{32}V_p^{23}X^1X^2\right)(V_y^{11} - V_y^{12}) \\ & - \left(H^{21}V_p^{12}X^1 + H^{22}V_p^{22}X^1 + H^{11}V_p^{11}X^2\right)(V_y^{22} - V_y^{23})X^3 < 0 \end{aligned}$$

As this determinant is non-zero, the Jacobian is non-singular and we can invoke the implicit function theorem to write $\boldsymbol{\theta}$ as a function of $\boldsymbol{g}$ such that $\boldsymbol{F}(\boldsymbol{\theta}^*(\boldsymbol{g}), \boldsymbol{g}) = \boldsymbol{0}$. (For clarity of notation going forward, I omit stars from variables at their equilibrium value.)

Differentiating the equilibrium condition with respect to g_2 yields

$$\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\theta}} \cdot \frac{\partial \boldsymbol{\theta}}{\partial g_2} = -\frac{\partial \boldsymbol{F}}{\partial g_2}$$

$$\begin{pmatrix} V_y^{11} - V_y^{12} & 0 & V_p^{11} & -V_p^{12} & 0 \\ 0 & V_y^{22} - V_y^{23} & 0 & V_p^{22} & -V_p^{23} \\ H^{11} & 0 & X^1 & 0 & 0 \\ -H^{21} & H^{22} & 0 & X^2 & 0 \\ 0 & -H^{32} & 0 & 0 & X^3 \end{pmatrix} \begin{pmatrix} \partial \tilde{y}_1 / \partial g_2 \\ \partial \tilde{y}_2 / \partial g_2 \\ \partial p_1 / \partial g_2 \\ \partial p_2 / \partial g_2 \\ \partial p_3 / \partial g_2 \end{pmatrix} = \begin{pmatrix} V_g^{12} \\ -V_g^{22} \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

Solving for $\partial\boldsymbol{\theta}/\partial g_2$ yields

$$\frac{\partial\boldsymbol{\theta}}{\partial g_2} = \left(\det\frac{\partial \boldsymbol{F}}{\partial\boldsymbol{\theta}}\right)^{-1} \cdot$$

$$\begin{bmatrix} -H^{22}X^1X^3V_p^{22}V_g^{22}\left(\frac{V_g^{12}}{V_g^{22}} - \frac{V_p^{12}}{V_p^{22}}\right) + V_g^{12}X^1X^2X^3\left(V_y^{22} - V_y^{23}\right) - H^{32}V_g^{12}V_p^{23}X^1X^2 \\ -H^{21}X^1X^3V_p^{22}V_g^{22}\left(\frac{V_g^{12}}{V_g^{22}} - \frac{V_p^{12}}{V_p^{22}}\right) - V_g^{22}X^1X^2X^3\left(V_y^{11} - V_y^{12}\right) + H^{11}V_g^{22}V_p^{11}X^2X^3 \\ H^{11}H^{22}X^3V_p^{22}V_g^{22}\left(\frac{V_g^{12}}{V_g^{22}} - \frac{V_p^{12}}{V_p^{22}}\right) - H^{11}V_g^{12}X^2X^3\left(V_y^{22} - V_y^{23}\right) + H^{11}H^{32}V_g^{12}V_p^{23}X^2 \\ -H^{11}H^{22}X^3V_g^{22}V_p^{11} - H^{21}H^{32}V_g^{12}V_p^{23}X^1 + H^{22}V_g^{22}X^1X^3\left(V_y^{11} - V_y^{12}\right) + H^{21}V_g^{12}X^1X^3\left(V_y^{22} - V_y^{23}\right) \\ -H^{21}H^{32}X^1V_p^{22}V_g^{22}\left(\frac{V_g^{12}}{V_g^{22}} - \frac{V_p^{12}}{V_p^{22}}\right) - H^{32}V_g^{22}X^1X^2\left(V_y^{11} - V_y^{12}\right) + H^{11}H^{32}V_g^{22}V_p^{11}X^2 \end{bmatrix}$$

The Jacobian determinant is negative and, using Lemmata 1 and 2, it is clear upon inspection that the signs of $\partial\boldsymbol{\theta}/\partial g_2$ are consistent with the statement of the proposition.

The proportionality of $\partial\tilde{y}_1/\partial g_2$ and $\partial p_1/\partial g_2$ follows from

$$\frac{-X^1}{H^{11}}\frac{\partial p_1}{\partial g_2} = \frac{\partial\tilde{y}_1}{\partial g_2}$$

□

Proposition 2.

$$\frac{\partial\tilde{y}_1}{\partial g_2} > 0$$

if and only if $M(\tilde{y}_2, p_2, g_2)$ is sufficiently large relative to $M(\tilde{y}_1, p_2, g_2)$ such that

$$\frac{M(\tilde{y}_2, p_2, g_2)}{M(\tilde{y}_1, p_2, g_2)} - 1 > \frac{H^3\dfrac{\partial V^3}{\partial\tilde{H}^3} + V_y(\tilde{y}_2, p_3, g_3) - V_y(\tilde{y}_2, p_2, g_2)}{H^2\dfrac{\partial V^2}{\partial\tilde{H}^2}} \tag{2.5}$$

where

$$\frac{\partial V^j}{\partial\tilde{H}^j} = \frac{V_p(\tilde{y}_2, p_j, g_j)}{\int_{\tilde{y}_j}^{\tilde{y}_{j+1}} h_p(p_j, y)f(y)\,\mathrm{d}y - S_p^j(p_j)} \tag{2.6}$$

and $H^j = h(p_j, \tilde{y}_2)f(\tilde{y}_2)$

Proof. Re-write $\partial\tilde{y}_1/\partial g_2$ as

$$\frac{\partial\tilde{y}_1}{\partial g_2} = \left(\det\frac{\partial\boldsymbol{F}}{\partial\boldsymbol{\theta}}\right)^{-1} X^1X^2X^3V_g^{12}\left[H^2\frac{\partial V^2}{\partial\tilde{H}^2}\frac{M^{22} - M^{21}}{M^{21}} + \left(V_y^{22} - V_y^{23}\right) - H^3\frac{\partial V^3}{\partial\tilde{H}^3}\right]$$

where $M^{ij} = M(\tilde{y}_i, p_j, g_j)$. Then

$$\begin{aligned}
&\frac{\partial \tilde{y}_1}{\partial g_2} > 0 \\
\Leftrightarrow\ & H^2 \frac{\partial V^2}{\partial \tilde{H}^2} \frac{M^{22} - M^{21}}{M^{21}} + \left(V_y^{22} - V_y^{23} \right) - H^3 \frac{\partial V^3}{\partial \tilde{H}^3} + V_y^{23} - V_y^{22} > 0 \\
\Leftrightarrow\ & \frac{M^{22} - M^{21}}{M^{21}} > \frac{H^3 \frac{\partial V^3}{\partial \tilde{H}^3} + V_y^{23} - V_y^{22}}{H^2 \frac{\partial V^2}{\partial \tilde{H}^2}}
\end{aligned}$$

□

Proposition 3. *The comparative statics for g_1 are g_3 are*

$$\frac{\partial \tilde{y}_1}{\partial g_1} > 0; \quad \frac{\partial \tilde{y}_2}{\partial g_1} > 0; \quad \frac{\partial p_1}{\partial g_1} > 0; \quad \frac{\partial p_2}{\partial g_1} < 0; \quad \frac{\partial p_3}{\partial g_1} < 0;$$

$$\frac{\partial \tilde{y}_1}{\partial g_3} < 0; \quad \frac{\partial \tilde{y}_2}{\partial g_3} < 0; \quad \frac{\partial p_1}{\partial g_3} < 0; \quad \frac{\partial p_2}{\partial g_3} < 0; \quad \frac{\partial p_3}{\partial g_3} > 0;$$

Proof. Following the proof of Proposition 1,

$$\frac{\partial \boldsymbol{\theta}}{\partial g_1} = \left(\det \frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\theta}} \right)^{-1} \cdot \begin{bmatrix} H^{22} V_g^{11} V_p^{22} X^1 X^3 + H^{32} V_g^{11} Y_p^{23} X^1 X^2 - V_g^{11} X^1 X^2 X^3 \left(V_y^{22} - V_y^{23} \right) \\ H^{21} V_g^{11} V_p^{22} X^1 X^3 \\ -H^{11} H^{22} V_g^{11} V_p^{22} X^3 - H^{11} H^{32} V_g^{11} V_p^{23} X^2 + H^{11} V_g^{11} X^2 X^3 \left(V_y^{22} - V_y^{23} \right) \\ H^{21} H^{32} V_g^{11} V_p^{23} X^1 - H^{21} V_g^{11} X^1 X^3 \left(V_y^{22} - V_y^{23} \right) \\ H^{21} H^{32} V_g^{11} V_p^{22} X^1 \end{bmatrix}$$

and

$$\frac{\partial \boldsymbol{\theta}}{\partial g_3} = \left(\det \frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\theta}} \right)^{-1} \cdot \begin{bmatrix} -H^{22} V_g^{23} V_p^{12} X^1 X^3 \\ -H^{11} V_g^{23} V_p^{11} X^2 X^3 - H^{21} V_g^{23} V_p^{12} X^1 X^3 + V_g^{23} X^1 X^2 X^3 \left(V_y^{11} - V_y^{12} \right) \\ H^{11} H^{22} V_g^{23} V_p^{12} X^3 \\ H^{11} H^{22} V_g^{23} V_p^{11} X^3 - H^{22} V_g^{23} X^1 X^3 \left(V_y^{11} - V_y^{12} \right) \\ -H^{11} H^{32} V_g^{23} V_p^{11} X^2 - H^{21} H^{32} V_g^{23} V_p^{12} X^1 + H^{32} V_g^{23} X^1 X^2 \left(V_y^{11} - V_y^{12} \right) \end{bmatrix}$$

Again following the methods used in Proposition 1, we get the sign of each element of these vectors. $\square$

Appendix C

Appendix to Chapter 3

C.1 Comparison with Competitive Price Ceiling Model

It appears that with lawyers, as with many professionals, their advice is self interested, to the possible detriment of their client. But it may be that this behavior is simply to avoid profit losses. For example if the bankruptcy market is competitive, the lawyers may not be able to make money on additional filings unless the price ceiling goes up. In that case the steering may not be due to poorly informed debtors being manipulated, but rather a natural result of the price caps imposed by the courts.

We think that a competitive model with zero profits on the margin runs into some difficulties that make it a less preferable explanation than the information and steering model. First, the price is not a true ceiling, since attorneys can petition for higher fees if needed. Second, in equilibrium judges are unlikely to leave fees low enough to price large numbers of people out of the market altogether, especially if they *want* people to file under Chapter 13. Beyond these institutional problems, the main problem is that in a price ceiling model, the elasticity we recover is the elasticity of supply, and a supply elasticity of approximately 0.3 is easiest to reconcile with the data using an information and steering model. We consider below the two best explanations for this elasticity that preserve a zero profits condition: capacity constraints and idiosyncratic cost or risk differences across debtors.

C.1.1 Capacity Constraints

A capacity constraint explanation for our estimate would be that for the market in an area to accommodate a 10% increase in quantity requires almost a 40% increase in price due to a sharply increasing marginal cost of production.[1] It would be odd if bankruptcy had such sharp supply constraints when similar activities don't seem to. Firms generally, and law firms specifically, are probably capable of filing 10% more paperwork without raising their prices at all, let alone 40%.

Fortunately, we have fairly strong evidence that price ceilings are not constraining firms. In September and October of 2005, when the Bankruptcy Abuse Prevention and Consumer Protection Act (BAPCPA) was about to be implemented, there was a huge demand shock as debtors raced to file pre-BAPCPA. While the average monthly filings of Chapter 13 in the summer of 2005 were between 30,000 and 35,000, in September there were 41,200 filings and in October there were 74,000 filings.[2] Under a binding price ceiling and capacity constraints, our estimated price elasticity would imply that the only way to get firms to have increased their supply so much would be for the price to go up ten-fold.[3] Since short-term elasticities are smaller than long-term ones, and supplying new filings should be even less elastic than switching between filings, this is a lower bound. A ten-fold increase seems implausible, as it should, because a supply elasticity of 0.3 driven by capacity constraints is probably implausible.

As a caveat, we should note that the 2005 law change was a huge shock to the national bankruptcy market, affecting both supply and demand. If this somehow fundamentally reordered the market, then perhaps the supply elasticity was vastly different in 2005 than

1. Moving 10 log points with an elasticity of 0.3 requires a 33 log point increase in price, which is a 39% increase.

2. These numbers are from the same bankruptcy data source used in the econometric work.

3. Doubling requires a log increase of .7. It would take a .7/.3 increase in price to do this—which is 2.33 log points above the current price.

in our sample year of 2007. We know that Chapter 13 bankruptcies rose at double digit rates from 2008 through 2010 fiscal years which would also be difficult to explain under the competitive model (U.S. Courts, 2013).[4]

C.1.2 Customer Heterogeneity

Instead of capacity constraints, firms may be constrained because idiosyncrasies of individual filers mean that they cannot make money on some filers at low prices. Thus supply curves slope up because the client population is changing. This may be because lawyers are moving to riskier borrowers, who are more likely to default or debtors' paperwork and filing expenses may become increasingly complex and costly.

The difficulty with this explanation is that, given competitive markets, the price charged will not be constant across debtors. Rather it will be a price schedule based on the debtor's expected riskiness and cost of paperwork. In this case, we should observe many agents at prices below the price ceiling. With our estimate of elasticity we can infer how much more expensive (in risk or costs) agents get as one moves through the distribution. Our estimate of 0.3 says that, on average, and over a wide range of the support, adding 10% on to the distribution of filers increases the marginal cost by 32 log points, which is, once again, almost a 40% cost increase. This massive difference across filers would be hard to explain with idiosyncratic changes in case complexity, although perhaps extremely high default rates could provide a pathway. But if this is so, the price distribution of Chapter 13 filings should be very disperse—with only 10% of the observed population in any given interval of a 40% increase in price. Yet we know that this is not the case. Figure 1 plots the density of the log Chapter 13 price and shows clearly that prices are heavily concentrated at the mean, presumably at the no-look fee.[5] What little dispersion we observe is an order of magnitude less than would be observed in the postulated market where many debtors are

4. With a binding price ceiling, demand shocks should not matter at all, and with a supply shock one can still back out from the change in quantity what the implied change in the cost curve must have been.

5. The results are the same when restricted to the top third of districts with the highest no-look fees where the price ceiling should be relaxed enough to reveal any latent differences across consumers.

vastly cheaper than others.

Moving outside a competitive setting, the effect of a price ceiling should go the opposite direction from what we found.[6] Given these reasons, we think the full information competitive model is not as easy to reconcile with the market as a model with information problems. Especially given that lawyers are explicitly hired and paid to provide information, so that an information model fits very naturally. In this case, the estimated elasticity suggests that it takes a 40% price increase to incentivize enough lawyers to switch 10% of filers between Chapters, which we think is plausible.

C.2 Supplementary Figures and Tables

6. Unless the price ceiling were so low that it drove the monopolist down to their marginal cost curve. But in that case they will look exactly like a competitive firm and the prior analysis applies.

Table C.1: *Wage Garnishments, Exemptions, and Price Data by State*

State	Protection from Garnishment	Garnishment Restriction Classification	Married Homestead Exemption	Married Homestead Exemption Classification	Average Chapter 13 Price Observations per District
Alabama	Federal	Light	$31,000	Low	12.5
Alaska	$602.50/wk	Heavy	$130,400	High	4.0
Arizona	Federal	Light	$115,000	High	4.0
Arkansas	Federal	Light	$31,350	Low	4.0
California	Federal	Light	$118,800	High	5.0
Colorado	Federal	Light	$70,500	Medium	4.0
Connecticut	$206/wk	Moderate	$224,000	High	4.0
Delaware	85% of "net income"	Moderate	$25,150	Low	4.0
District of Columbia	Federal	Light	$61,800	Medium	6.0
Florida	100% of D.I.	Heavy	Unlimited	High	5.5
Georgia	Federal	Light	$17,400	Low	42.0
Hawaii	80% of income	Moderate	$83,150	Medium	4.0
Idaho	Federal	Light	$72,600	Medium	4.0
Illinois	85% of D.I.	Moderate	$22,900	Low	15.0
Indiana	Federal	Light	$20,000	Low	5.0
Iowa	$206/wk	Moderate	Unlimited	High	4.0
Kansas	Federal	Light	Unlimited	High	4.0
Kentucky	Federal	Light	$27,600	Low	4.5
Louisiana	Federal	Light	$111,000	Medium	4.7
Maine	$226/wk	Moderate	$49,800	Medium	4.0
Maryland	Federal	Light	$17,000	Low	4.0
Massachusetts	Federal	Light	$129,500	High	4.0
Michigan	$206/wk	Moderate	$61,800	Medium	4.5
Minnesota	$206/wk	Moderate	$463,100	High	4.0
Mississippi	Federal	Light	$170,500	High	4.0

Table C.1: *(continued)*

Missouri	90% of income	Moderate	$25,800	Low	5.5
Montana	Federal	Light	$140,000	High	4.0
Nebraska	85% of D.I.	Moderate	$37,800	Medium	4.0
Nevada	Federal	Light	$277,000	High	4.0
New Hampshire	$257/wk (a)	Heavy	$89,600	Medium	4.0
New Jersey	90% of D.I.	Moderate	$61,800	Medium	12.0
New Mexico	$206/wk	Moderate	$112,500	Medium	4.0
New York	90% of D.I.	Moderate	$36,000	Medium	4.8
North Carolina	75% of D.I. (b)	Heavy	$32,500	Low	8.0
North Dakota	$206/wk	Moderate	$172,400	High	4.0
Ohio	Federal	Light	$18,100	Low	16.5
Oklahoma	Federal	Light	Unlimited	High	4.0
Oregon	$206/wk	Moderate	$54,000	Medium	4.0
Pennsylvania	100% of income	Heavy	$61,800	Medium	5.0
Rhode Island	Federal	Light	$203,700	High	4.0
South Carolina	100% of income	Heavy	$19,900	Low	8.0
South Dakota	80% of D.I. (b)	Heavy	Unlimited	High	4.0
Tennessee	Federal	Light	$29,300	Low	13.0
Texas	100% of income	Heavy	Unlimited	High	14.3
Utah	Federal	Light	$56,000	Medium	4.0
Vermont	85% of D.I. (b)	Heavy	$171,800	High	6.0
Virginia	Federal	Light	$54,000	Medium	5.0
Washington	$294/wk	Moderate	$105,400	Medium	5.5
West Virginia	80% of D.I.	Moderate	$37,400	Medium	4.0
Wisconsin	80% of D.I.	Moderate	$109,400	Medium	4.0
Wyoming	Federal	Light	$49,800	Medium	4.0

Notes: D.I. is disposable income. States that either forbid garnishment or limit garnishment to exclude "living expenses", which are marked with (b), are coded as heavy because this typically allowed debtors to evade substantive wage garnishment. (a) Garnishments are only effective for one paycheck at which point a new lawsuit must be filed. (b) "Living expenses" are exempt.

Table C.2: *Full Results from Table 3.3, Column 3*

Log 13/7 Fee Difference	0.32***	Age Below 6	0.44
	[0.082]		[0.907]
Log Repayment Rate	0.33**	Age 6 to 18	0.68
	[0.160]		[0.436]
Log Monthly Income	-0.17	Age 19 to 24	-1.01**
	[0.123]		[0.405]
Log(Secured Debt/Total Debt)	-0.07	Age 25 to 29	0.29
	[0.221]		[0.511]
Urban	-0.25***	Age 30 to 39	-0.97**
	[0.052]		[0.431]
Population/1000	-0.13	Age 40 to 49	-0.63
	[0.145]		[0.442]
Population/1000 Sq.	0.08	Age 50 to 59	-0.64
	[0.247]		[0.391]
Married	-0.82**	Unemployed	0.06
	[0.376]		[0.717]
Divorced	-0.77*	Self-Employed	-1.42***
	[0.399]		[0.350]
Household of 2	-0.97**	Household Income under $10,000	-2.73***
	[0.368]		[0.678]
Household of 3	-0.54	Household Income $10-$20,000	-3.07***
	[0.591]		[0.434]
Household of 4	-1.31***	Household Income $20-$30,000	-1.64***
	[0.448]		[0.432]
Household of 5	-0.5	Household Income $30-$40,000	-1.27***
	[0.474]		[0.397]
Household of 6	-0.22	Household Income $40-$50,000	-0.82**
	[0.729]		[0.343]
Household over 6	-0.78	Household Income $50-$60,000	-0.99**
	[0.722]		[0.385]

Table C.2: *(continued)*

Female Head of Household	0.04	Household Income $60-$75,000	-0.52
	[0.291]		[0.367]
Head of Household Below Age 24	0.77	Household Income $75-$100,000	-0.37
	[0.673]		[0.337]
Finished High School	-0.47*	Fraction Homeowners	0.47**
	[0.267]		[0.186]
Finished college	-0.44*	25th Pctle. of Log Housing Value	-0.14***
	[0.217]		[0.050]
Black	0.74***	75th Pctle. of Log Housing Value	-0.16**
	[0.071]		[0.075]
Hispanic	-0.16		
	[0.204]		
Other Race	-0.26		
	[0.239]		

Notes: Dependent variable is Log Fraction 13. OLS Regression includes state fixed effects. Standard errors in brackets are clustered at the state level. * significant at 10%, ** significant at 5%, *** significant at 1%. Each zip code in the regressions is weighted by the number of bankruptcies that occurred there in 2007.