

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Essays in Environmental and Labor Economics

Permalink

<https://escholarship.org/uc/item/1vg6d8xj>

ISBN

9798186381587

Author

Bian, Luorao

Publication Date

2026-06-06

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Essays in Environmental and Labor Economics

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Luorao Bian

Committee in charge:

Professor Olivier Deschênes, Chair
Professor Tamma Carleton, University of California, Berkeley
Professor Heather Royer

June 2026

The Dissertation of Luorao Bian is approved.

Professor Tamma Carleton, University of California, Berkeley

Professor Heather Royer

Professor Olivier Deschênes, Committee Chair

May 2026

Essays in Environmental and Labor Economics

Copyright © 2026

by

Luorao Bian

Acknowledgements

Life is all about experiences, and the best part is exploring the new and making sense of the unknown.

I am deeply thankful to Jeffrey Grogger, who first introduced me to the field and modeled what it means to be an economist. Working with Jeff helped me develop a scholarly perspective that has become the bedrock of my career.

I am profoundly grateful to Olivier for his generous guidance and continuous support over the years. His unique perspective helped me navigate the field of environmental economics and build a strong interdisciplinary research portfolio.

I wish to express my sincere thanks to Tamma for her invaluable advice and insightful feedback. From Tamma, I learned how to approach complex questions and tackle challenges, which significantly elevated the quality of my work.

I am incredibly grateful to Heather Royer, whose mentorship shaped not only my research but also my professional trajectory. From Heather, I learned how to effectively advocate for my work and myself as a scholar.

My sincere thanks also go to Gonzalo, Peter, Ted, my fellow graduate students, and the dedicated faculty and staff in the Economics Department for their ongoing support.

Finally, I am eternally grateful to my parents, to Shu-Chen, and to my friends and family. Thank you for encouraging me to pursue this path and for standing by me every step of the way.

Curriculum Vitæ

Luorao Bian

Education

2026	Ph.D. in Economics (Expected), University of California, Santa Barbara.
2021	M.A. in Economics, University of California, Santa Barbara.
2019	M.A. in Public Policy, The University of Chicago.
2018	B.A. in Public Affairs Management and Finance (double major), Central University of Finance and Economics.

Fields of Study

Environmental, Labor, and Urban Economics

Publications

Bian, L., 2025. Do microclimates matter? The Health Impacts of Urban Heat Using Fine-scale Data. *Journal of Environmental Economics and Management*, 133, 103191.

Bastian, J., **Bian, L.** and Grogger, J., 2022. How Did Safety-Net Reform Affect the Education of Adolescents from Low-Income Families?. *Labour Economics*, 77, 102031.

Bastian, J., **Bian, L.** and Grogger, J., 2021. How Did Safety-Net Reform Affect Early Adulthood among Adolescents from Low-Income Families?. *National Tax Journal*, 74(3).

Abstract

Essays in Environmental and Labor Economics

by

Luorao Bian

This dissertation consists of three chapters. The first chapter examines the health impacts of urban heat and its implications for adaptation. Cities tend to be warmer than their suburban counterparts, a phenomenon known as the heat island effect. This effect varies significantly over time and across different neighborhoods. Using high-resolution, satellite-derived surface temperature data and emergency medical service records from Northern California, this chapter employs fixed-effects Poisson regressions to analyze these variations. The findings show that neighborhood-level extreme heat has significant adverse health effects. However, urban heat—a highly localized and predominantly anthropogenic component of overall temperature—is less harmful than temperature shifts on a larger geographic scale. This is likely because people can better adapt to urban heat. These results highlight the importance of considering both overall temperature exposure and adaptation potential in urban environments. The effects of urban heat also vary by season; while additional warmth can provide protection during colder periods, it poses substantial health risks on warmer days. Finally, policy simulations suggest that mitigating extreme urban heat could effectively reduce health risks, particularly during warm periods. Tree planting, especially in less vegetated areas, emerges as a beneficial strategy for protecting urban populations.

The second chapter studies the unintended environmental externalities and spillovers of place-based policies, which broadly aim to boost local economies and improve welfare. Specifically, it examines city minimum wage ordinances, a popular place-based labor

regulation intended to raise earnings and living standards for low-wage workers in over 60 U.S. cities. Using variation in timing and intensity across 17 California cities in the 2010s, I show that higher local wage floors increase rents and incentivize workers to live and work in neighboring areas outside the policy jurisdiction. These spatial shifts alter commuting patterns, increase traffic congestion, and worsen air quality, with neighboring areas more severely affected than policy cities, despite not being the direct targets. This worsening of air quality disproportionately affects poor and minority communities due to suboptimal amenities and limited adaptation capacity, rather than worker re-sorting. Simulations combining a traffic-emissions model with a particle-dispersion model attribute up to 2.8% of traffic-related emissions and 15 additional annual deaths to the ordinances. For low-wage workers, the costs of these unintended externalities offset 40% of the intended wage gains. In summary, the results suggest that environmental externalities, though indirect, can shift the welfare impacts of place-based policies far from their intended goals.

The final chapter, co-authored with Jeffrey Grogger and Jacob Bastian, focuses on the long-run impacts of U.S. safety net reform. Roughly 25 years ago, the U.S. safety net was substantially overhauled. Here, we ask how those reforms affected the educational attainment of youths who were teenagers at the time of implementation. We take a difference-in-differences approach, following adolescents from two generations roughly 20 years apart. In each generation, we compare two groups: one more likely to have been affected by the reforms, and one less likely. Under some assumptions, our approach identifies the joint, or bundled, effects of the constituent policy changes that make up safety net reform. We find evidence that the reforms may have reduced educational attainment for women, while having small positive effects on education for men. Finally, we offer suggestions as to why our findings differ from those of previous studies examining individual components of safety net reform.

Contents

Curriculum Vitae	vi
Abstract	vii
1 Do Microclimates Matter? The Health Impacts of Urban Heat Using Fine-Scale Data	1
1.1 Introduction	1
1.2 Heat Island Effect	4
1.3 Data and Preliminary Analysis	8
1.4 Empirical Strategy	15
1.5 Results	19
1.6 Urban Heat Mitigation	33
1.7 Conclusion	39
2 The Environmental Implications of Place-Based Policy: Evidence from Minimum Wage Ordinances	41
2.1 Introduction	41
2.2 Background	46
2.3 Data and Summary Statistics	52
2.4 Empirical Strategy	61
2.5 Results	64
2.6 Heterogeneous Impacts	73
2.7 Policy Simulation	83
2.8 Conclusion	91
3 How Did Safety-Net Reform Affect the Education of Adolescents from Low-Income Families?	93
3.1 Introduction	93
3.2 Background	97
3.3 Data	103
3.4 Estimation and Results	106

3.5 Discussion	129
A Appendix for Chapter 1	140
B Appendix for Chapter 2	160
C Appendix for Chapter 3	173
D Permissions and Attributions	177

Chapter 1

Do Microclimates Matter? The Health Impacts of Urban Heat Using Fine-Scale Data

1.1 Introduction

Urban residents typically experience warmer climates than their suburban counterparts, a phenomenon known as the heat island effect. In the United States, approximately 80% of the population resides in urban areas, where temperatures are, on average, $3.7^{\circ}F$ higher than in suburban areas (*Census Bureau, 2020*, Hsu et al. 2021). For comparison, global warming—another significant anthropogenic temperature shift—has increased global temperatures by $1.9^{\circ}F$ since the pre-industrial era (Hansen et al. 2010).

The heat island effect varies substantially across seasons and neighborhoods, resulting in highly localized temperature fluctuations. To examine its health implications, I use fine-resolution data from Northern California to first determine whether health outcomes change with temperature variations at the neighborhood level. I then explore the an-

thropogenic nature of the heat island effect and demonstrate how its impact differs from those of regional climate fluctuations, which I interpret as arising from differences in adaptation potential. Finally, I quantify the total health impact of the heat island effect and use policy simulations to provide guidance on effective public health interventions.

Using surface temperature data from NASA’s Moderate Resolution Imaging Spectroradiometer (MODIS) and local emergency medical service (EMS) records from San Francisco, San Jose, San Rafael, and other cities in Marin County, I employ reduced-form fixed-effect Poisson regressions to estimate the impact of daily temperature fluctuations on morbidity at the census-block-group level. I then follow Rizwan et al. (2008), Imhoff et al. (2010), Schwarz et al. (2011), Zhou et al. (2013), and Manoli et al. (2020) to decompose overall temperature into background temperature (the average temperature in the urban buffer, constant for the entire city) and urban heat (local deviations from the urban buffer temperature). I show that prior studies implicitly assume equal health impacts for both components at any given overall temperature level, despite the fact that urban heat is highly localized and more sensitive to human modifications. Their differing adaptation potentials explain my empirical finding that urban heat, a micro-level temperature shift, is relatively less harmful than macro-level variations (background temperature).

I further employ similar Poisson regressions to estimate the total health impact of urban heat. The benchmark regressions are estimated separately by season because urban heat essentially alters temperature relative to some baselines while the health-temperature relationship is nonlinear, and because the two temperature components exhibit seasonal patterns. My approach incorporates macroclimate conditions nonparametrically through fixed effects and implicitly models the interactions between the two temperature components as part of seasonal heterogeneity. Finally, I compare the effectiveness of various urban heat mitigation strategies through simulations, including tree planting—one of the most widely adopted urban cooling policies in U.S. cities.

The first key finding is that local heat spikes increase morbidity risk at the neighborhood level. More importantly, urban heat—the locally varying component—poses relatively lower health risks compared to background temperature. I interpret this result as people can better adapt to urban heat. This study makes two primary contributions. First, I use novel datasets to recover the health-temperature relationship at the census-block-group level, offering insights unavailable in studies that analyze data at the county level or larger (Deschênes & Moretti 2009, Deschênes & Greenstone 2011, Barreca et al. 2016, Carleton 2017, Mullins & White 2019, and Carleton et al. 2022). Second, by decomposing overall temperature into two components with different adaptation potentials, I implicitly approximate the adaptation margin. The results highlight the importance of considering not only overall temperature exposure but also how effectively people can adapt. Even if total heat exposure remains constant, health risks can differ depending on adaptation effectiveness and feasibility. These findings also suggest that the common practice of assuming spatial separability in damage functions may not produce the most precise estimates in urban settings (Hsiang 2016).

The second finding offers a novel assessment of urban heat’s health impacts in different seasons. The heat island effect exacerbates health risks during warm periods but mitigates them on cold days. This pattern is driven by both the nonlinear health-temperature relationship and seasonal interactions between the two temperature components. Further evidence suggests that the spatial and temporal distribution of urban heat is more influential in shaping health outcomes than demographic and socioeconomic factors. These analyses have three key strengths. First, I model the heat island effect dynamically rather than treating it as a static, location-specific attribute (e.g., imperviousness). The use of high-frequency, day-to-day variation exploits the temporal dynamics of the heat island effect, which is less likely to correlate with income and other confounders that remain relatively stable over time. Second, even after accounting for regional macroclimate con-

ditions and location-specific characteristics, I find that neighborhood-level temperature variation significantly affects health outcomes in one of the wealthiest American communities in the late 2010s. While a higher income does not fully offset climate impacts. Third, unlike studies focusing on specific extreme events or short time windows, my analysis spans the entire year, allowing for a more comprehensive assessment of urban heat impacts and policy benefits. My estimates reveal not only the adverse effects of urban heat in warm months (Tan et al. 2010, Heaviside et al. 2016, Heaviside et al. 2017, and Chakma et al. 2024) but also its protective effects in winter (Macintyre et al. 2021). These findings highlight the nonlinear health-temperature relationship and seasonal dynamics in shaping the health impacts of heat island effect.

In terms of policy application, I simulate the benefits of urban heat mitigation and find that reducing extreme levels of urban heat alone could lower emergency medical caseloads by up to 2.63% on hot days—achieving nearly half the reduction that would result from completely eliminating urban heat. Finally, I examine the health benefits of tree planting, one of the most common urban heat mitigation strategies.¹ I find that tree planting is almost always beneficial, particularly when prioritizing less green neighborhoods.

1.2 Heat Island Effect

The heat island effect plays a crucial role in shaping neighborhood-level climates. This phenomenon refers to urban areas being significantly warmer than their surroundings, primarily due to the enhanced heat retention of urban environments. The dense infrastructure of buildings, roads, pavements, and other urban elements absorbs and emits more heat compared to natural landscapes such as grasslands, forests, and bodies of

¹Examples of tree planting programs include the California Urban Forestry Program, the New York State Urban Forestry Grants, and the Urban Forestry Program of the District of Columbia.

water. These factors contribute to higher temperatures in urban settings. Additionally, human activities, such as vehicle operation and industrial production, generate excess heat. Urban geometries further exacerbate the effect by reducing wind speed, which impedes heat dispersion from buildings and slows the natural cooling process.


The heat island effect is measured by calculating the temperature difference between urban areas and the average temperature in the urban buffer. For each pixel j in city c on day d of year y , the temperature reading represents the overall temperature L , which is decomposed into background temperature T and urban heat UH . Specifically, I draw on the works of Tan et al. (2010), Clinton & Gong (2013), Zhou et al. (2013), and Quan et al. (2014) to decompose overall temperature using Equation 1.1:

$$L_{jcdy} = UH_{jcdy} + T_{cdy} \quad \text{where; } T_{cdy} = \frac{1}{N_c} \sum_{k \in Buffer} L_{kcdy} \quad (1.1)$$


Urban areas are identified as pixels classified as "Urban and Built-up Lands" under the International Geosphere-Biosphere Programme (IGBP), which includes areas with "*at least 30% impervious surface area, including building materials, asphalt, and vehicles.*" I construct a 5-km urban buffer around urban pixels, excluding adjacent built-up pixels. Next, I calculate the average temperature across all buffer pixels, which serves as the background temperature for all pixels j in city c on day d of year y . In other words, all pixels within the same city share the same T on a given day. Finally, I determine urban heat by computing the difference between the overall temperature in each pixel and the background temperature T . Figure 1.1 illustrates the division of urban areas and urban buffers, along with the average level of urban heat over the entire sample period. The choice of buffer remains consistent over time and has minimal impact on the estimation after accounting for fixed effects, as demonstrated in Table A.1.

Figure 1.1: Average Urban Heat by Urban Area and Buffer Division

(a) Urban Areas, Buffers, and Average Urban Heat



(b) Average Urban Heat at the Pixel Level in a San Jose Neighborhood



Notes: The inner circles outline urban clusters defined by the MODIS MCD12 Land Type product, while the outer circles indicate a 5-km buffer zone around these clusters. The urban buffer refers to the area between the urban clusters and the buffer boundary. Urban heat is measured as the deviation of local daily temperature from the buffer average. For pixels within the buffer, urban heat is nonzero if the local temperature differs from the buffer average. Some areas near the buffer are excluded due to missing land cover data.

The decomposition of overall temperature divides total heat into two parts with distinct characteristics. Background temperature T is defined as the buffer average, representing regional climatic conditions over a relatively large area, while urban heat is highly localized, capturing how much local climate conditions deviate from the regional average. In terms of origin, background temperature T is primarily driven by natural weather fluctuations, whereas urban heat is largely anthropogenic, influenced by human activities that modify the local energy budget.² As a result, urban heat is more responsive to human adaptation and mitigation efforts.

These distinctions suggest differences in adaptation potential, which may ultimately lead to divergent health impacts of the two temperature components. The conceptual framework in Appendix A.0.1 shows that the total health impact of temperature can be attributed to both biological and adaptation channels. Urban heat is highly localized and largely human-made, offering a greater range of adaptation possibilities and facilitating more effective adaptation. To illustrate this, holding overall temperature L constant and all else equal, a high UH (and thus a low T) indicates that a heat spike is local, making surrounding areas more likely to remain cool. Temperature differences promote faster heat dissipation by Fourier’s Law, while relocation to cooler neighborhoods is more feasible and less costly. In contrast, a high T suggests a low UH given a fixed L , meaning that heat spikes are more likely to be regional. In such cases, the aforementioned mitigation strategies are less effective since other areas are more likely to share the same degree of heat. As a result, ventilation is less effective, and relocating to cooler neighborhoods becomes more difficult or even impossible. Given the significant benefits of adaptation, as documented in Deschênes (2014), Barreca et al. (2015), Barreca et al. (2016), and Carleton et al. (2022), differences in adaptation potential may lead to different total heat impacts.

²An energy budget accounts for the balance of energy entering and exiting a system.

These considerations are often overlooked in existing approaches that estimate the health-temperature relationship, which account only for the total amount of heat but not its composition (Deschênes & Moretti 2009, Mullins & White 2019, Carleton et al. 2022). In Section 1.4.2, I propose a set of empirical tests to examine whether these differences matter in an urban setting and outline an alternative estimation framework to address the issue in Section 1.4.3.

From a policy perspective, differences in adaptation potential suggest that mitigating the heat island effect can be more feasible to protect people from climate change. While practitioners may have limited ability to alter regional climate patterns, they can still invest in heat island mitigation measures—such as tree planting, installing cool roofs, and adding reflective pavement—to improve local microclimates. Additionally, policymakers can facilitate adaptation by setting up heat alerts to help people avoid localized heat hotspots.

1.3 Data and Preliminary Analysis

In this section, I first introduce the construction of health and temperature variables. Next, summary statistics and seasonal patterns are presented to describe the data, inform the empirical strategy, and facilitate the interpretation of empirical results in Section 1.5.

1.3.1 Data

To estimate the health risks associated with microclimate changes at a fine scale, I assembled a novel dataset comprising emergency medical service (EMS) records from the cities of San Francisco, San Jose, and Marin County, along with NASA satellite images covering these areas. The analysis sample includes areas that received EMS services in San Francisco, San Jose, San Rafael, Novato, and other smaller cities in Marin County.

Health is measured in terms of acute morbidity, defined as the total number of EMS transfers at the daily census-block-group level. EMS transfers refer to incidents where patients receive immediate medical attention and require extended medical observation. The raw data is obtained from local EMS administrations and fire departments, containing both the geographical coordinates and time of incidents. I match these coordinates with the 2010 census block group shapefile and further aggregate the data to the census-block-group-by-date level.³ For further quality assurance, I conduct robustness checks using EMS incidents that require ambulance dispatches, and find even more significant results. I also compare my measure with California vital records on deaths, which explains 79.0% of the variation in death counts, with a correlation coefficient of 0.890.⁴

The main temperature measures are overall temperature L and its two components, urban heat UH and background temperature T , all at the block-group-by-date level, obtained from the Land Surface Temperature and Emissivity swath (MOD21A1D). The raw data provides daily images of daytime land surface temperature at a 1-km resolution. Following Imhoff et al. (2010), Schwarz et al. (2011), Manoli et al. (2019), and other studies on the heat island effect, I use surface temperature to measure urban heat, taking advantage of richer climate information at a fine (1-km) geographical scale (Jin & Dickinson 2010). As discussed in the previous section, temperature decomposition is conducted at the grid-cell level to maintain precision. The grid-level measures are then aggregated to the block-group-by-date level by matching grid cells with census block group shapefiles, using area weighting following Carleton (2017). That is, for each block group, I

³Patient home location data is unavailable. Using incident location may introduce endogeneity concerns, as people may actively avoid locations that are unfavorable for health. In that case, the actual heat exposure experienced by the patient may be greater than my observation, thus biasing the estimates toward zero.

⁴While morbidity and mortality are highly correlated and both increase in response to temperature spikes, on cold days, morbidity decreases while mortality increases (Gould et al. 2024). Consequently, one should not expect a perfect correlation between these two health measures, nor infer mortality impacts on cold days from this study.

calculate the overlapping area with the temperature grid and compute an area-weighted average temperature. The use of population weights makes minimal difference, as census block groups are small, whereas population density data is coarser in geographical scale. Additional details about the process are summarized in the data appendix.

An alternative temperature measure is inverse-distance-weighted air temperature collected from meteorological stations. However, such data typically provides only a few data points per city (e.g., three in San Francisco, two in San Rafael, and three in San Jose) and fails to capture potentially drastic temperature variation within urban areas. Chakma et al. (2024) show that such a measure can even be misleading. In contrast, satellite data captures temperature variations at the neighborhood level, providing a reasonable approximation of heat exposure. Studies suggest that although numerical differences exist between surface temperature and station data (e.g., higher variance and mean for surface temperature), they exhibit high correlations, follow similar temporal patterns, and yield comparable conclusions in economic studies (Benali et al. 2012, Heft-Neal et al. 2017).

For a quality check, I show in the Data Appendix that the temperature-morbidity relationship based on EMS caseloads and surface temperature exhibits similar patterns to those estimated from hospital emergency department visits and station-derived temperature in Gould et al. (2024).

1.3.2 Summary Statistics and Seasonal Patterns

As discussed in Section 1.3, the following analyses use surface temperature and the number of EMS transfer records at the census-block-group-by-date level. The sample is restricted to areas within the 2010 *Census Urban Areas* and within the jurisdiction of EMS providers in San Francisco, San Jose, San Rafael, and the rest of the urban

areas in Marin County (Appendix Figure A.1).⁵ The sample spans from January 2015 to December 2019.

Table 1.1 presents summary statistics for temperature and health variables. Given the diversity of health and climate conditions across cities, I also report city-specific statistics. Panel A reports the mean, standard deviation, and the 10th, 25th, 75th, and 90th percentiles for three temperature measures. For the entire sample, the average surface temperature is $81.765^{\circ}F$ with a standard deviation of $17.763^{\circ}F$. Both are relatively high, given that they are based on satellite-derived surface temperatures rather than conventional air temperature. Nonetheless, surface temperature allows for a better characterization of urban climate dynamics and has minimal impact on the interpretation of the results. Regarding the two temperature components, background temperature averages $76.509^{\circ}F$, while the heat island effect averages $5.261^{\circ}F$. This indicates that, on average, urban areas are $5.261^{\circ}F$ warmer than the buffer average. Notably, the 10th percentile of urban heat is negative ($-2.126^{\circ}F$), suggesting that it is not uncommon for some urban census block groups to be cooler than the buffer average.

⁵Because *Census Urban Areas* are defined based on more sophisticated criteria beyond *Urban and Built-up Lands*, areas outside the urban buffer may also be included, accounting for $< 1\%$ of the sample.

Table 1.1: Summary Statistics

Panel A: Temperatures		Share of variations by sources							
		Mean	S.D.	10th	25th	75th	90th	Geographical	Seasonal
A1: The entire sample, N=1033081									
Overall temperature L	81.765	17.763	58.109	66.513	95.774	106.287	0.097	0.746	0.001
Background temperature T	76.509	15.897	54.978	62.342	89.556	97.414	0.016	0.850	0.004
Urban heat UH	5.261	6.146	-2.126	1.512	9.292	13.160	0.401	0.156	0.005
A2: San Rafael and other cities in Marin, N=232431									
Overall temperature L	75.630	15.306	55.164	62.342	87.384	95.634	0.056	0.761	0.003
Background temperature T	74.364	14.074	54.915	61.446	86.531	91.590	0.000	0.863	0.004
Urban heat UH	1.271	5.270	-5.224	-1.649	4.370	7.988	0.464	0.030	0.001
A3: San Francisco, N=238020									
Overall temperature L	77.010	14.023	57.866	65.030	87.734	94.442	0.036	0.758	0.003
Background temperature T	73.826	14.417	54.654	60.702	86.338	91.989	0.000	0.864	0.005
Urban heat UH	3.179	6.222	-4.634	-0.036	6.899	10.573	0.189	0.280	0.006
A4: San Jose, N=562630									
Overall temperature L	86.310	18.804	59.987	69.170	102.110	109.854	0.014	0.878	0.001
Background temperature T	78.530	16.889	55.305	63.439	93.429	100.576	0.000	0.885	0.004
Urban heat UH	7.790	5.168	1.624	4.353	11.335	14.689	0.173	0.436	0.012
Panel B: N EMS cases									
B1: The entire sample, N=1033081	Mean	S.D.	Minimum	10th	25th	75th	90th	Maximum	Share of $N \geq 1$
B1: The entire sample, N=1033081	0.337	0.703	0.000	0.000	0.000	0.000	1.000	8.000	0.248
B2: Cities in Marin, N=232431	0.220	0.494	0.000	0.000	0.000	0.000	1.000	4.000	0.188
B3: San Francisco, N=238020	0.383	0.830	0.000	0.000	0.000	1.000	1.000	8.000	0.255
B4: San Jose, N=562630	0.366	0.713	0.000	0.000	0.000	1.000	1.000	8.000	0.270


Notes: Land surface temperature has larger variances and means, but it strongly correlate with air temperature (Jin & Dickinson 2010, Benali et al. 2012, and Heft-Neal et al. 2017). The last three columns report the share of variations by three sources, obtained from a three-way ANOVA. Neighborhoods with the same urban buffer share the same background temperature, so there's no geographical variation in T in A2-A4.

Additionally, I conduct a three-way ANOVA to decompose variation by time trend (calendar year), geography (block group), and season (day of the year), presented in the last three columns. This exercise serves two purposes: to elucidate the factors contributing to different types of heat and to ensure sufficient residual variation remains (presumably random) after controlling for fixed effects. For overall temperature L , most of the variation (74.6%) comes from seasonality, with geography and calendar year trends contributing 9.7% and 0.1%, respectively. Similarly, background temperature variation is predominantly driven by seasonality. In contrast, 40.1% of urban heat variation is attributed to geography, indicating that neighborhood characteristics play a significant role. Seasonality contributes 15.6%, suggesting that day-to-day fluctuations in urban heat are common and not solely dictated by neighborhood differences, even though socioeconomic conditions are correlated with differential urban heat exposure (Benz & Burney 2021, Borsky et al. 2024). A substantial portion of variation remains unexplained, which is consistent with the fact that urban heat is sensitive to factors such as wind, cloud cover, and human activities. Panels A2 to A4 present the same statistics disaggregated by city, demonstrating that the level, dispersion, and sources of variation in urban heat are city-specific.


Seasonal patterns are illustrated in Figure 1.2. The top panel plots the average levels of L , T , and UH by day of the year. By construction, the average UH is the area between L and T , which varies considerably across seasons. The bottom panel presents city-specific averages of UH and T by month.

Figure 1.2: Seasonality of Overall Temperature, Background Temperature, and Urban Heat

(a) Daily Averages Across the Entire Sample



(b) Monthly Averages by City



Notes: Panel A presents the average overall temperature T and land surface temperature L for each day of the year, with urban heat UH represented as the difference between the two. Panel B shows average monthly T and UH by sub-sample.

Both panels align with Zhou et al. (2013), who also identify hysteretic cycles between urban and background temperatures—i.e., background temperature T initially increases more slowly than urban temperature but catches up later.

These seasonal patterns have implications for the empirical strategy that estimate the impacts of UH, as they suggest seasonal dependence between microclimate and macroclimate measures. Details of my empirical strategy are discussed in Section 1.4.3.

1.4 Empirical Strategy

In this section, I describe the empirical approaches that first identify the health impact of neighborhood-level temperature variation L , then, test for the potential differential impacts of urban heat UH versus background temperature T , and finally, measure the casual effects of urban heat.

1.4.1 Health Impacts of Temperature Variations at Neighborhood Level

To connect to the literature, I first estimate the health impact using actual variation in the data. In other words, I test whether health responds to overall temperature L at a fine geographical scale using high-frequency local temperature variation. I adapt from Deschênes & Moretti (2009) to employ the following fixed-effect Poisson regressions, where temperature is categorized into bins to account for nonlinear temperature effects and accommodate count data:

$$\log E(y_{icdy}) = \sum_k \beta_k \times \mathbb{I}\{L_{icdy} \text{ in bin } k\} + \epsilon_i + \epsilon_d + \epsilon_y + \epsilon_{icdy} \quad (1.2)$$

In this equation, for each census block group i in city c on day d of year y , the expected number of emergency cases follows a Poisson process as a function of 14 five-degree L bins, with the 6th bin $L \in (75, 80]$ omitted as the reference group. The coefficients of interest, β_k , represent the proportional change in expected caseload in response to changes in the

overall temperature bin, relative to the reference level.

The identification assumption is that, conditional on the fixed effects, the set of climate measures is uncorrelated with the error term, consistent with Deschênes & Moretti (2009), Barreca (2012), Barreca et al. (2016), Carleton (2017), and White (2017). In an urban setting, one might worry about confounding factors as the literature has shown that demographic and socioeconomic conditions correlate with heat exposure (Benz & Burney 2021, Borsky et al. 2024, Chakma et al. 2024). However, this issue is less likely to bias the results given the use of high-frequency data. By leveraging daily variations conditional on location-specific effects, it is less plausible that income, race, education, and other confounders systematically covary with daily climate shocks. For robustness checks, Appendix Table A.2 and Figures A.2 and A.4 show that the results hold even with more flexible fixed effects (block-group-by-month and block-group-by-week). One might still be worried that other climate conditions such as humidity may vary across neighborhoods and confound the results. In this case, denser urban areas tend to be drier (*IPCC 2021, EPA 2023*), and lower humidity facilitates human thermoregulation under extreme heat. Therefore, the estimated health impact is likely a lower bound of the effect that would be observed under constant humidity conditions. Results on the health impact of neighborhood-level temperature variation are presented in Section 1.5.1.

1.4.2 Health Impacts by Temperature Component

This section assesses the differential health impacts of the two heat components as Section 1.2 highlights potential distinctions between urban heat UH and background temperature T. The highly localized and largely human-made nature of urban heat suggests greater adaptation potential, which may lead to differential health impacts.

I estimate the following specification, modifying Equation 1.2 by adding some controls

for urban heat:

$$\log E(y_{icdy}) = \sum_k \beta_k \times \mathbb{I}_k\{L_{icdy} \text{ in bin } k\} + \gamma UH_control + \epsilon_i + \epsilon_d + \epsilon_y + \epsilon_{icdy} \quad (1.3)$$

The empirical approach relies on the following intuition: if urban heat UH and background temperature T contribute similarly to health risks, then overall temperature L should be sufficient to capture the temperature effects, while adding explicit controls for UH should have little impact on the estimates of L. Thus, γ should be approximately zero. Since changes in L arise from either UH or T, controlling for UH implies that variation must come from T, making the estimate of L conditional on UH effectively represent the impact of T. Here, *UH_control* is broadly defined as some function of urban heat, and I test multiple specifications to confirm that functional form assumptions do not drive the results.

Proposition 2 in Appendix A.0.1 details the interpretation of the results. In summary, changes in biological and/or adaptation channel may influence total health impacts, following Grossman (1972), Becker & Murphy (1988), Moretti & Neidell (2011), Barreca et al. (2016), and Deschênes et al. (2017). If differential health impacts are observed between UH and T, they stem from differences in biological response, adaptation, or both. Given the nature of UH and T as discussed in Section 1.2, more effective adaptation to urban heat is a more plausible explanation. The results are discussed in Section 1.5.2.

1.4.3 Health Impacts of Heat Island Effect

To estimate the casual impacts of urban heat, I estimate the following fixed-effect Poisson regression:

$$\log E(y_{icdy}) = \sum_p \rho_p \times \mathbb{1}_{\{UH_{icdy} \text{ in bin } p\}} + \epsilon_i + \epsilon_{cdy} + \epsilon_{icdy} \quad (1.4)$$

The benchmark estimates of the heat island effect are presented on a seasonal basis, as the health-temperature relationship is nonlinear (Deschênes & Moretti 2009, Deschênes & Greenstone 2011, Mullins & White 2019, Carleton et al. 2022). The presence of UH effectively adds additional degrees to certain temperature ranges. Imposing a constant effect across seasons implies that an additional degree of UH has the same health risk on cold and hot days, contradicting existing knowledge. This approach also accounts for the seasonal patterns of UH and T as in Section 1.3.2 by allowing for estimates to vary across seasons while avoiding making explicit assumptions on how they interact.

Equation 1.4 follows the structure of Equation 1.2, with two key differences. First, the parameters of interest, ρ_p , indicate proportional changes in expected caseload due to changes in UH bins (11 two-degree bins, with $UH \in (-1, 1]$ omitted as the reference group). Second, instead of separate day-of-the-year and calendar-year fixed effects, I use ϵ_{cdy} , a city-by-date fixed effect that accounts for background temperature, humidity, pressure, and other macroclimatic conditions in a non-parametric manner. The use of high-frequency data conditional on location-specific effects further mitigates concerns that demographic and socioeconomic factors drive the estimates.

Seasonal groupings are determined via cross-validation. I begin with 10 different seasonal groupings, each covering 2–4 months with different start and end points. For each alternative, I use 5-fold cross-validation to compute the out-of-bag MSE in the health-UH regression and select the grouping with the lowest MSE. The main results are


presented in Section 1.5.3, with Appendix Figure A.5 confirming robustness to alternative groupings based on T. Section 1.5.4 further explores the heterogeneity by the distribution of UH as well as demographic and socioeconomic conditions.

1.5 Results

1.5.1 Impacts of Overall Temperature

Figure 1.3 reports the extent to which overall temperature affects health across urban neighborhoods using observed temperature variation from the data, L. As discussed in Section 1.4, the estimates represent proportional changes relative to the reference group. For example, the > 115 bin has a coefficient of 0.0668, indicating a 6.68% increase in caseloads relative to the reference bin. This suggests that even at a fine geographical scale, higher overall temperature L increases health risks.

Figure 1.3: Contemporaneous Effects of Overall Temperature L on EMS Caseload



Notes: This figure presents the estimated relative change in EMS caseload (y-axis) associated with different L bins (x-axis). Estimates are derived from Equation 1.2, controlling for day-of-the-year, calendar-year, and block-group fixed effects. Confidence intervals are at the 95% level, with standard errors clustered at the census block level. The reference category is $L \in [75, 80)$ (the red $\times$).


The results are robust to alternative outcome measures and model specifications, as shown in Appendix Figure A.2. As a placebo test, I also include overall temperature on the following day and find little impact, as shown in Appendix Figure A.9.

1.5.2 Differential Health Risks

This section compares whether the health impact of urban heat differs from that of background temperature across multiple specifications. Figure 1.4 compares the estimates of L bins in Figure 1.3 to a set of estimates further conditioned on linear UH (Equation 1.3 with continuous UH). Per the discussion in Appendix A.0.1, larger estimates (in absolute value) in this specification indicate greater health impacts when the change in heat comes entirely from background temperature T, i.e., $\frac{dH}{dL} > \frac{dH}{dL}|\overline{UH}$, or equivalently, $\frac{dH}{dT} > \frac{dH}{dUH}$. Equivalently, UH estimate (γ in Equation 1.3) is -.0022 with a p-value of less than 0.001. Negative UH estimate suggests that, holding the overall

temperature bin fixed, more heat coming from urban heat is relatively better for health. UH and L estimate taking opposite signs suggest health is more sensitive to changes in T than to changes in UH.

Figure 1.4: Contemporaneous effects of Overall Temperature L on EMS Caseload, by UH Controls




Notes: The figure above compares the changes in EMS caseload (y-axis) due to changing L bins (x-axis), with and without holding a constant UH. Estimates are obtained from two regressions. The blue (round) estimates are specified by Equation 1.2 as in the previous figure 1.3, while the orange (square) ones are the same estimates but further conditional on constant UH, i.e., adding continuous UH into Equation 1.2. They are both estimated with day-of-the-year, calendar-year, and block-group fixed effects. The estimate of continuous UH in the second regression is -0.0022 with P-value 0.000, indicating that holding overall temperature bin fixed, more heat coming from urban heat is relatively good for health. Confidence intervals are at the 95% level, with S.E. clustered at the census block level. $L \in [75, 80)$ (the red $\times$) is omitted as the reference group.

To show that the differential impacts of UH and T are not driven by functional form, Figure 1.5 further demonstrates that health impacts vary by heat composition even within the same L bin. This highlights that we should consider not only the level of overall temperature but also its composition. For this exercise, I re-estimate the health impacts given in Figure 1.3 by high versus low urban heat subgroups defined within each 5-degree L bin. Now, urban heat controls are more flexible and can vary across L bins. That is, I further break down the 14 overall temperature indicators in Equation 1.2 into high and low urban heat subgroups, producing 28 groups. High versus low urban heat is defined in two ways. In Figure 1.5, I hold overall temperature "constant" (within 1-degree L) and then group observations by urban heat percentile. This method imposes a strict


restriction on overall temperature L but allows for geographical and temporal variations. Within each L bin, the group with a lower level of urban heat UH (and thus a higher level of background temperature T , as the total of the two is fixed) is more severely impacted, indicating that, holding overall temperature fixed, having more heat from urban heat UH is better than having more from background temperature T . For a causal interpretation, it assumes that, conditional on (almost) the same L and fixed effects, high versus low UH is as good as random, given 1-degree L and the set of fixed effects. An alternative comparison is conducted in Figure 1.5, in which the grouping is conditional on the same census block group within the same 5-degree L bin and the same month. In other words, the comparison is intertemporal within the same location and month, across days with higher versus lower levels of urban heat. Regardless of the method, both figures support the previous finding that heat composition matters and that urban heat UH is relatively less harmful than background temperature T . Estimates are robust to more flexible model specifications, including controlling for block-group-by-month or block-group-by-week fixed effects, or changing the outcome variable to the number of dispatches at the same level, as shown in Appendix Figure A.3.

Figure 1.5: Effects of L by High-Low UH Group

(a) Cross-Block-Group Comparison



(b) Within-Block-Group Comparison



Notes: Both panels compare the health impacts of L by high versus low levels of UH. For each 5-degree L bin (as in Equation 1.2), observations are further divided into high versus low levels of UH by interacting L bin indicators with high versus low UH indicators. The top panel compares high-low UH within the same 1-degree bins, i.e., comparing observations of the same L under high and low UH conditions. Similarly, the bottom panel conducts the same comparison but with an alternative UH definition. Here, UH is defined within the same census block group, the same 5-degree L bin, and the same month. Intuitively, it compares the effects of L within the same location and season on days with high versus low UH. The reference group consists of L between 75-80 degrees with relatively low urban heat. Estimates control for day-of-the-year, calendar-year, and block-group fixed effects. 95% confidence intervals are reported, with standard errors clustered at the census block level. The P-value tests the null hypothesis that the effects of L are the same for all high versus low urban heat groups.

I employ even more flexible models by first allowing temperature effects to vary by season and to put UH in bins, as shown in Appendix Table A.2. As discussed in Section 1.4.2 and Appendix A.0.1, nonzero urban heat estimates indicate that, holding overall temperature fixed, health impacts differ by heat composition and the effects of UH differ from those of T. Conversely, zero UH estimates imply that the effects of UH are similar to the effects of T but not that UH has null influence on health. In addition, I estimate the same health impacts using cross-validated random forest classifiers with overall temperature L, urban heat UH, census block group, and date as predictors. Adding urban heat UH reduces the prediction MSE by 4.11%, compared to the model without UH. In the specification with both L and UH, UH serves as the most important feature in predicting health outcomes, with a feature importance of 0.319.


Thus far, multiple pieces of evidence have demonstrated that the health impact of urban heat differs from that of background temperature. Combining these findings with the discussion in Section 1.2 and Appendix A.0.1, I interpret the results as evidence that people can more effectively adapt to urban heat than to background temperature.

1.5.3 Impacts of Heat Island Effect

Figure 1.6 reports the effects of urban heat in May–July and October–February using 12 two-degree UH bins. Seasonal variations in the UH effect arise primarily from the nonlinearity of the health-temperature relationship, as discussed in Section 1.4.3. That is, UH estimation essentially models how health outcomes change with additional heat at different parts of the temperature distribution. A nonlinear relationship implies that the effects of the same UH level can vary depending on whether the baseline temperature is low, moderate, or high. Furthermore, since urban heat estimates capture both the direct health impacts of urban heat and its interaction with macroclimatic conditions,

UH effects may vary even within the same temperature window.

Figure 1.6: Effects of Urban Heat on EMS Caseload, by Season



Notes: Estimates are based on Equation 1.4, with **city-by-date**, calendar-year, and block-group fixed effects. Background temperature T is also at the city-by-date level and is controlled non-parametrically by the fixed effect. There are 11 UH bins, with the 4th bin, $UH \in [-1, 1)$ (the red x), omitted as the reference group. The estimates indicate the relative changes in caseload due to changing UH bins. For instance, the first UH bar in the top panel indicates that in May–July, moving from $UH \in [-1, 1)$ to $UH < -5$ leads to a 3.8% increase in caseloads. Within this bin and season, the average levels of T and L are 89 and 80 degrees, respectively. Estimates for the rest of the months (March–April and August–September) are omitted as they are close to zero.

In May–July (the top panel), a low level of urban heat is associated with a moderate increase in caseloads. More importantly, when urban heat exceeds five degrees, the increase in caseloads becomes significant (up to 9.35% more EMS cases relative to the reference bin). During this period, background temperature is relatively high, and health outcomes are highly sensitive to small changes in L . Adding additional urban heat increases, L and consequently, the caseload, as shown in the middle-to-right portion of Figure 1.5. By contrast, on days with extremely low UH during this period, most heat comes from T . Since health is more sensitive to changes in T than in UH, having most

of the heat come from T results in a net positive effect of urban heat. However, in October–February (the bottom panel), a higher level of urban heat appears to be beneficial for health (reducing EMS cases by up to 6.53%), as days with higher UH during colder periods experience relatively lower health risks. Effects in the remaining months are not significant and are therefore omitted.

The results are robust to grouping by T (Appendix Figure A.5), using an alternative outcome measure, and including block-group-week or block-group-month fixed effects (Appendix Figure A.4). Placebo tests examining whether urban heat tomorrow affects health today show little impact (Appendix Figure A.10).


1.5.4 Heterogeneous Impacts of Heat Island Effect

In this subsection, I first demonstrate that the distribution of heat across geographical and temporal dimensions plays a critical role in characterizing the health-urban-heat relationship. I then present evidence that demographic and socioeconomic factors, such as income, exert less influence.


I begin by examining whether the cross-sectional geographical distribution of urban heat matters. This analysis is motivated by the findings in Section 1.5.2, which indicate that health is more sensitive to changes in large-scale events, such as background temperature T , than to changes in UH at a smaller scale. Figure 1.7 shows that health is more sensitive to urban heat when heat stress is more uniformly distributed across neighborhoods (the low-variance group). The differences in health impacts between the low- and high-variance groups are statistically significant from October to February ($p = 0.005$) but not in May to July ($p = 0.110$).

Figure 1.7: Effects of Urban Heat by the Variability of Heat

(a) Same-Day Comparison: By Geographical Variation



(b) Same-Place Comparison: By Historical Variation



Notes: Estimates are based on Equation 1.4, with UH bins interacting with high-low $Var(UH)$ indicators. The top panel defines $Var(UH)$ within the same city and same day, dividing block groups into two zones based on geographical location. For each day (same T), $Var(UH)$ is calculated for each zone, labeled as high or low, and then interacted with UH bins. This compares the effects of UH between zones with more evenly distributed (low $Var(UH)$) and more dispersed (high $Var(UH)$) urban heat. The bottom panel calculates variance within the block group across days, highlighting the learning from prior exposure to diverse urban heat conditions. The P-value tests the null hypothesis that the effects of UH are the same for high versus low variance groups.


To construct these groups, I disaggregate each urban heat bin in Equation 1.4 based on the cross-sectional geographical dispersion of urban heat. Specifically, I first use geographical coordinates and clustering method to divide census block groups within each city into two zones. For each zone on a given day, I compute the cross-sectional variance of urban heat and classify the zones as either high- or low-variance groups. This approach ensures the the comparison is within the same city on the same date to eliminate confounders such as background temperature T and macroclimatic conditions. A lower UH variance indicates a more homogeneous heat distribution within each zone. The findings suggest that highly localized heat spikes are less influential, possible due to adaptation being more effective feasible as in Section 1.2. In contrast, geographical sorting, an alternative explanation, appears less convincing given that the split is approximately 40-60 between two zones, and differences in exposure level and annual household income are minimal (-0.38 degrees and \$716.90, respectively).

Since prior exposure may influence individuals' ability to cope, I next examine whether people who experience more variable UH levels are less affected. The hypothesis is that greater uncertainty in UH may lead to better preparedness. I measure uncertainty using the variance of urban heat within each block group over time. I define high- versus low-variance groups conditional on city. The findings indicate that individuals exposed to more variable UH are less affected by extreme UH in May–July ($p = 0.010$). Additional comparisons show that differential responses cannot be explained by sorting based on either exposure (a difference of -0.0389 degrees) or annual household income (a difference of -\$1569.451). These results support the adaptation hypothesis and remain robust to alternative fixed effects, as shown in Appendix Figures A.7 and A.6. I also test whether prior exposure levels shape responses. The group with lower prior exposure is more responsive to moderate increases in UH in May–July than the high-exposure group, but the differences are not statistically significant ($p = 0.199$; see Appendix Figure A.8).

To address concerns regarding geographical sorting, I assess whether urban heat exposure levels differ and whether people differentially respond to the same urban heat shocks by income and demographic factors. Table 1.2 reports differences in heat exposure level across quartiles of demographic and socioeconomic characteristics, based on American Community Survey 5-Year Data (2011-2015). I focus on household income, the share of the nonwhite population, the share of the Hispanic population, and the share of the population aged 65+, as prior studies have documented disparities in heat exposure and health responses by race, income, and demographics (Benz & Burney 2021, Borsky et al. 2024). The results indicate that temperature differences across demographic and economic quartiles are modest. Block groups with a higher share of minority and/or Hispanic populations are, on average, exposed to higher levels of UH and L (2.183 and 2.213 degrees, respectively). Conversely, areas with a higher share of older populations and/or higher household income experience less severe exposure (-1.513 degrees in UH and -1.625 in L). However, temperature variability remains similar across groups. Overall, these findings align with prior research showing that minority and lower-income populations face moderately higher heat stress (Benz & Burney 2021, Borsky et al. 2024).

Figure 1.8 compares the health responds differentially to the same UH shocks for high- and low-income groups. The sample is split into high-income (highest and higher quartiles) and low-income (lower and lowest quartiles) categories and I test whether the same UH shock affects health by the same proportion. High UH imposes greater increases in caseload for the low-income group in May–July, whereas the high-income group experiences smaller increases. However, the difference is not statistically significant ($p = 0.368$). Both groups respond more similarly in October–February, with ($p = 0.815$). Similar comparisons for other demographic and socioeconomic factors in Appendix Figures B.3 to B.4 also fail to reject the null hypothesis that neighborhoods respond differently to the same urban heat based on demographic or socioeconomic status.

Figure 1.8: Limited Differential Impacts of UH on EMS Caseload by Income Group



Notes: Estimates control for city-by-date and block-group fixed effects, with UH bins interacting with two income indicators. Within each panel, results are organized by rows (seasons) and by high versus low proportions of the demographic group. The P-value tests the null hypothesis that the effects of UH are the same across high and low income groups.

Table 1.2: Urban Heat Exposure by Demographics and Socioeconomic Status

	L Mean	L S.D.	UH Mean	UH S.D.	N EMS Mean	N EMS S.D.	N
% Nonwhite							
4th quartile	82.457	17.662	5.897	5.673	0.355	0.673	262,950
3rd quartile	81.765	17.659	5.272	6.035	0.399	0.807	256,769
2nd quartile	81.450	17.849	4.960	6.319	0.365	0.760	257,554
1st quartile	81.370	17.861	4.900	6.487	0.226	0.528	255,808
% Hispanic							
4th quartile	82.838	17.901	6.302	5.860	0.479	0.824	262,642
3rd quartile	82.073	17.815	5.582	5.912	0.366	0.732	258,122
2nd quartile	81.490	17.669	5.008	6.162	0.248	0.561	257,308
1st quartile	80.625	17.587	4.119	6.434	0.250	0.638	255,009
% population aged 65+							
4th quartile	80.651	17.655	4.265	6.230	0.306	0.646	261,493
3rd quartile	81.676	17.788	5.176	6.214	0.277	0.625	258,258
2nd quartile	82.396	17.795	5.875	6.009	0.341	0.706	258,811
1st quartile	82.357	17.759	5.747	5.992	0.424	0.815	254,519
Median household earning							
4th quartile	80.947	17.532	4.476	6.158	0.265	0.665	262,210
3rd quartile	81.744	17.763	5.265	6.039	0.211	0.485	257,329
2nd quartile	81.817	17.894	5.312	6.301	0.355	0.717	258,102
1st quartile	82.572	17.828	6.011	5.984	0.518	0.857	255,440

Notes: Quartiles are calculated separately for three city clusters (San Jose, San Francisco, and San Rafael, and other cities in Marin County). Summary statistics using 5-year ACS data at census block group level.

1.6 Urban Heat Mitigation

Previous sections indicate that the health impacts of urban heat are influenced by its spatial and temporal distribution, with adaptation playing a pivotal role. In this section, I conduct three exercises to analyze how urban heat mitigation can protect health.

First, I use empirical estimates from Section 1.5.3 to simulate proportional changes in health outcomes under different urban heat mitigation scenarios. Second, I examine the importance of explicitly accounting for urban heat in estimation by comparing simulations generated using two estimation frameworks: one that follows a traditional approach without distinguishing heat composition (Equation 1.2) and another that explicitly accounts for urban heat (Equation 1.4). Finally, I connect these simulations to actual policy variables to provide insights into effective urban heat mitigation strategies that prioritize health. Additional details on the simulation process are provided in Appendix A.0.2.


To compare the benefits of various urban heat mitigation strategies, I calculate predicted caseloads under mitigated UH scenarios versus actual UH. Let $\hat{y}_b$ denote the predicted caseload with actual urban heat and $\hat{y}_h$ the caseload under mitigated UH. I compute the predicted proportional change in caseload as $\frac{\hat{y}_h - \hat{y}_b}{\hat{y}_b}$ for all observations and report the average. I consider four urban heat reduction scenarios: (1) a uniform 4-degree reduction in UH; (2) a marginal reduction such that each block group achieves the same UH level as the coolest neighboring block group on a given day; (3) targeting extreme UH levels by trimming the distribution to the range $[-5, 5]$; and (4) completely eliminating UH.⁶ The fourth scenario provides an approximation of the total health impact of the heat island effect in this region, serving as a benchmark for comparison.

⁶The "coolest neighbor" refers to the block group within the same census tract with the lowest UH level on a given day. For instance, if three block groups in the same tract have UH levels of 3.2, 6.7, and 9.2 degrees, the counterfactual assumes all three block groups adopt the lowest level, 3.2 degrees. This assumption is reasonable because neighboring block groups are more likely to share similar climatic and socioeconomic conditions.

Figure 1.9 reports the proportional change in caseload between the actual and simulated cases at different levels of background temperature T . When background temperature is low, reducing UH is slightly counterproductive, as urban heat provides some insulation against cold. Eliminating UH in such conditions increases the caseload by up to 0.91%. Other mitigation strategies yield similar but marginally negative effects. As background temperature rises, the benefits of UH mitigation become more pronounced. Even minor adjustments, such as mimicking the UH level of the coolest neighbor, lead to a 0.51% reduction in EMS incidents (with an average UH reduction of 1.3°F). A uniform 4-degree reduction provides much larger benefits, reducing caseloads by up to 2.02% when T is high. However, targeting extreme UH (with an average reduction of 2.3°F) yields even greater benefits, reducing caseloads by up to 2.63%. The most substantial benefits come from eliminating UH entirely, with caseload reductions reaching 5.25% on the hottest days. Comparing the first three strategies to total UH elimination, their effectiveness ranges from 10% to 50% of the full elimination benchmark, indicating that these mitigation approaches can significantly reduce the health burden of the urban heat island effect.

The bottom panel of Figure 1.9 presents the distribution of 2-degree urban heat across different background temperature levels. Darker red indicates higher UH, and a larger red area corresponds to a greater proportion of extreme UH occurrences. The shape of the red area closely follows the curvature of the four lines in the top panel, suggesting that the benefits of urban heat mitigation largely depend on how much extreme UH is reduced. These results indicate that targeting extreme urban heat is the most effective strategy.

Figure 1.9: Effects of UH Mitigation on EMS Caseload by Hypothetical Mitigation Methods




Notes: This analysis considers four urban heat (UH) mitigation strategies: (1) a uniform reduction of 4 degrees in UH, (2) a marginal improvement where each block group adopts the UH level of its coolest neighboring block each day, (3) a targeted reduction that trims extreme UH levels to fall within $[-5, 5]$, and (4) complete elimination of UH. The top panel presents the proportional reduction in EMS caseload as a function of background temperature. The comparison is based on predicted caseloads under actual UH versus reduced UH, as estimated using Equation 1.4. The bottom panel illustrates the distribution of UH across temperature levels. As before, UH is categorized into 2-degree bins, with red shades indicating a higher likelihood of high UH and blue shades indicating a higher likelihood of low UH.

To assess the importance of explicitly accounting for UH in estimation, I compare the marginal health impacts across estimation models. Figure 1.10 contrasts the previous simulations (Figure 1.9) with results estimated from Equation 1.4, which relies only on overall temperature and ignores heat composition. This comparison highlights the extent to which estimation methods influence predicted health benefits. In cold periods, the model relying on overall temperature L suggests UH mitigation is beneficial, contradicting the UH model. While in warm seasons, the model relying on L shows that the benefits of

UH mitigation kick in when background temperature reaches 85 degrees and the benefits get smaller when T is between 95 to 105 degrees. In contrast, the model that accounts for UH explicitly suggests that UH mitigation contributes to health at 75 degrees while the benefits steadily increase until T reaches 105. In addition, the latter model also predicts larger health benefits of urban heat mitigation in warm seasons.

Finally, recognizing that the scenarios above may not be fully achievable through real-world policies, I examine how tree-planting can enhance health outcomes across seasons by increasing tree canopy. This simulation follows a two-step process. First, I estimate the effects of trees on urban heat using machine learning models and vary tree canopy levels to obtain UH associated with certain levels of tree canopy. Second, I predict caseloads based on these UH levels and calculate the proportional change in caseload. Additional details on this methodology are provided in the Appendix A.0.2.

Figure 1.10: Differential Urban Heat Mitigation Benefits by Estimation Methods




Notes: This figure compares the estimated benefits of UH mitigation based on two different approaches: (1) using overall temperature (L) alone and (2) explicitly accounting for urban heat (UH). The orange (short-dashed) line represents results from previous figures, where the same reduction in UH was estimated using overall temperature (L) only. These estimates suggest that lower UH reduces health risks, with mitigation benefits becoming evident at relatively high background temperatures. In contrast, the UH-based estimates suggest that while reducing UH may not be beneficial in colder periods, the benefits of UH mitigation begin at much lower background temperatures.

Figure 1.11 demonstrates that tree planting effectively promotes health but exhibits diminishing marginal returns. When tree canopy is increased uniformly across all block groups, the first 5-percentage-point increase leads to a substantial reduction in EMS cases, whereas subsequent increases of 5 or 10 percentage points generate smaller marginal benefits. An alternative policy approach targets neighborhoods with the lowest tree canopy levels, ensuring that no neighborhood falls below a minimum threshold. Despite requiring a smaller overall increase in tree canopy, this method produces effects comparable to or greater than the uniform increase. Given that the mean tree canopy coverage is 0.113

(the national average is 0.271), the average canopy increases for the targeted approach are +0.039, +0.078, and +0.120 (row 2 in Figure 1.11), which are smaller than the +0.05, +0.10, and +0.20 increments in the uniform approach (row 1). These results suggest that prioritizing low-canopy areas—typically associated with lower-income and/or higher minority populations—can maximize the health benefits of urban heat mitigation.⁷ The bottom panel of Figure 1.11 presents the average UH reduction associated with each level of tree canopy modification. Trees function as a temperature regulator, protecting people from extreme heat on hot days.

⁷Behavioral responses also play a role; see Borsky et al. (2024) for further discussion.

Figure 1.11: Effects of Increasing Tree Canopy on EMS Caseload



Notes: The top panel shows the reduction in EMS caseload as a function of background temperature, comparing predicted caseloads under actual UH versus UH reduced through tree canopy expansion. Similar to previous simulations, this analysis is based on Equation 1.4, with the relationship between UH and tree canopy estimated separately using a random forest model (details in the data appendix). Two tree planting strategies are considered: (1) a uniform increase in tree canopy across all block groups (blue lines) and (2) a targeted increase in tree canopy for block groups with the lowest initial canopy cover (gray lines). For each method, varying treatment intensities are examined, and their impact on predicted caseloads is reported. The mean tree canopy cover in this sample is 0.113, compared to the U.S. average of 0.271 (Nowak et al. 2001).

1.7 Conclusion

Both global warming and the heat island effect are anthropogenic temperature increases, though they operate at different scales. While the impacts of global warming are well-documented, the heat island effect remains less understood, despite strong evidence linking extreme heat to adverse health outcomes at the regional level.

This study examines the health implications of neighborhood-level temperature variation and the role of the heat island effect using high-resolution data from Northern California. By integrating satellite-derived surface temperature with local emergency medical service (EMS) records, I quantify the morbidity impact of localized temperature shifts.

The initial findings indicate that localized heat spikes significantly increase morbidity risk at the neighborhood level. However, urban heat—a more localized temperature fluctuation—is relatively less harmful than broader regional temperature changes. This distinction underscores the importance of not only measuring overall temperature exposure but also considering how effectively people can adapt when assessing urban health impacts.

The second key finding provides a novel assessment of urban heat’s health effects across seasons. Results show that localized heat spikes can either exacerbate or mitigate health risks, depending on background temperature. Moreover, cross-sectional and historical variations in urban heat play a more critical role than demographic and socioeconomic factors, such as income.

Finally, I simulate the benefits of urban heat mitigation, showing that reducing extreme heat alone could eliminate nearly half of EMS cases attributed to the heat island effect. I also assess the health benefits of tree planting, a common urban heat mitigation strategy, and find that it is consistently beneficial—especially when prioritizing less green neighborhoods.

Chapter 2

The Environmental Implications of Place-Based Policy: Evidence from Minimum Wage Ordinances

2.1 Introduction

Place-based policies have been widely implemented to enhance economic performance in defined geographic areas and improve overall welfare. One well-known example is the Empowerment Zone Program, which provides tax incentives to attract businesses, increase employment, and reduce poverty (Neumark & Kolko 2010; Busso et al. 2013). However, the impacts of these policies are multifaceted, extending beyond employment to include effects on the housing market, the environment, and public health (Rossi-Hansberg et al. 2010; Grossman 2019; Koster & Van Ommeren 2019; Hua et al. 2023; Fan et al. 2023). Among these, policies such as local minimum wage ordinances target disadvantaged populations in specific geographical areas, classifying them as place-based strategies focused on people (Ladd 1994). In the U.S., more than 60 cities have adopted

city-level ordinances on top of state or federal minimum wage mandates. The policy goal is to ensure fair compensation and improve living standards for low-wage workers within the jurisdiction (Dube & Lindner 2021; Dube & Lindner 2024).

Place-based policies, by definition, aim to influence economic activity in specific geographic areas (Neumark & Simpson 2015). In general equilibrium, people and firms are mobile, and such policies may lead to the spatial displacement of both people and businesses (Neumark & Nizalova 2007; Faggio 2019; Duranton & Venables 2020). In the case of local minimum wage ordinances, the impacts are broad, spanning wages and employment as well as prices (e.g., restaurant menus, groceries, rents) and worker relocation (Katz & Krueger 1992; Card & Krueger 1994; Dube et al. 2007; Cadena 2014; MaCurdy 2015; Dube et al. 2016; McKinnish 2017; Cengiz et al. 2019; Leung 2021; Dustmann et al. 2022; Minton & Wheaton 2022; Borg et al. 2022). Changes in the spatial distribution of people affect traffic patterns, vehicle emissions, and, consequently, air quality. For urban residents, traffic congestion is a major quality-of-life issue (Anderson 2017). Pollutants such as nitrogen oxides (NO_2) and fine particulate matter ($PM_{2.5}$) from vehicle operations also contribute to air pollution, harming both human health and the environment (Currie & Walker 2011; Chen & Whalley 2012; Zhong et al. 2017).¹² Impacts on air quality and traffic congestion may therefore generate welfare losses and undermine the policy’s intended goals.

To identify the congestion and environmental impacts of local minimum wage ordinances, I exploit variation in the timing and levels of 17 city-level ordinances in Cal-

¹ $PM_{2.5}$ refers to fine inhalable particulate matter with a diameter of 2.5 micrometers or less. It can penetrate deep into the lungs and even enter the bloodstream. Exposure to $PM_{2.5}$ has been linked to irregular heartbeat, aggravated asthma, reduced lung function, and premature death, particularly among individuals with heart or lung disease (Turner et al. 2011; Kim et al. 2015; Feng et al. 2016; Burnett et al. 2018; Anderson 2020; Liang et al. 2020).

²Nitrogen dioxide (NO_2) is a gaseous pollutant and one of six criteria pollutants regulated by national air quality standards. NO_2 exposure has been associated with harmful respiratory diseases, including reduced lung function, increased asthma attacks, and airway irritation (Chauhan et al. 2003; Hesterberg et al. 2009; Hamra et al. 2015; Crouse et al. 2015).

ifornia.³ My identification strategy leverages treatment intensity (higher versus lower levels of local minimum wages) to capture impacts from both initial implementation and subsequent adjustments. I quantify the resulting traffic and air quality consequences by combining real-time freeway monitoring data with 1-km resolution remote-sensing data. My results show broad unintended impacts, ranging from vehicle miles traveled, vehicle hours traveled, and congestion probability to ambient concentrations of $PM_{2.5}$ and NO_2 . To explain the mechanism, I distinguish between workers' residential and workplace locations, analyze how ordinances affect the spatial distribution of workers across both margins, and demonstrate the implications for commuting patterns. I then introduce a conceptual framework that incorporates housing market equilibrium to identify key factors moderating residential re-sorting and provide empirical evidence on distributional impacts. To isolate traffic-related air quality impacts, I implement a two-stage simulation: estimating pollutant emissions from on-road traffic and tracing how these emissions translate into ambient air quality and mortality across neighborhoods using the Emission FACTors (EMFAC) model and the Intervention Model for Air Pollution (InMAP).

I first find that city-level minimum wage ordinances exacerbate traffic congestion and worsen air quality, particularly in areas adjacent to policy cities—even though these areas are not directly targeted by the policies. I document higher probabilities of congestion in both policy and neighboring areas, with the latter experiencing more vehicle miles traveled, longer vehicle hours, and higher concentrations of both $PM_{2.5}$ and NO_2 . These patterns are consistent with shifts in the spatial distribution of workers away from policy cities to surrounding areas, as higher minimum wages raise rents more sharply in policy cities. Diverging living costs incentivize workers to reside farther from their jobs and commute longer distances to avoid higher rents while still earning a higher wage.

³Policy cities include San Diego, Los Angeles, San Jose, Oakland, San Mateo, Berkeley, and 11 other cities, as well as areas neighboring these cities in California. See Table 2.1 for more details.

A second set of results shows that air quality impacts are disproportionately concentrated in neighborhoods with higher shares of Hispanic, service-sector, or less-educated workers. Part of this pattern reflects marginally disproportionate traffic impacts. More importantly, these neighborhoods possess fewer environmental amenities that mitigate pollution exposure, as they are closer to highways and lack green space. In other words, differences in adaptation capacity further explain the larger divergences in air quality impacts. By contrast, the evidence suggests that workers are not sorting into these neighborhoods, even though Hispanic, service-sector, and less-educated workers are significantly less likely to remain in policy cities. In terms of residential re-sorting, the conceptual framework predicts that workers value amenities (e.g., access to business establishments) and commute time. Empirically, amenity-rich neighborhoods are indeed more attractive to workers. There are some divergences with respect to travel time, but these are statistically insignificant.

Finally, to quantify the total contribution of minimum wage ordinances to traffic-related emissions and ambient air quality, I simulate emissions and concentrations under counterfactual scenarios without city-level ordinances. I find that removing these ordinances generates effects comparable to a 10% change in gasoline prices. Overall, ordinances contribute up to 2.8% of vehicle emissions, leading to 0.46% higher traffic-related pollutant concentrations, an additional 2.0 $PM_{2.5}$ -related deaths, and 13.2 traffic fatalities per year. Importantly, air quality impacts are three times larger in areas that were already highly polluted prior to the ordinances, meaning these neighborhoods bear the bulk of the health risks. A back-of-the-envelope calculation suggests that congestion and air quality impacts substantially erode wage gains, offsetting up to 40% of the benefits for low-wage workers.

This paper makes three contributions. First, it contributes to the literature on place-based policy. The findings not only reveal the displacement impacts of these policies but

also quantify the congestion and environmental costs in the U.S. setting.⁴ This novel evidence demonstrates that the indirect impacts of place-based policies can be costly and can substantially undermine their goals. The findings also highlight that negative policy externalities can be more severe in “untreated” neighboring areas than in targeted regions, emphasizing the importance of accounting for externalities and displacement effects in policy design (Neumark & Simpson 2015; Moretti 2024). More specifically, the paper contributes to the minimum wage literature. By explicitly distinguishing between residence and workplace locations, the analysis provides fine-resolution evidence of spatial displacement, discusses how it differs across margins, and reveals its implications for commuting patterns. These findings complement studies that consider only residence or workplace locations in state-level policy changes at a coarser geographic scale (Cadena 2014; Martin & Termos 2015; McKinnish 2017; Shirley 2018; Monras 2019; Dustmann et al. 2022).⁵

Second, it contributes to the literature on residential income sorting. The analysis shows that Hispanic, less-educated, and lower-wage workers are more affected, with residential re-sorting primarily driven by amenities (e.g., commute access, proximity to businesses). This provides empirical evidence on amenity-based sorting and income heterogeneity in urban residential equilibrium (Glaeser 2008; Gaigné et al. 2022). Moreover, it demonstrates that wage shocks at the bottom of the income distribution can alter prices and displace low-income households to suburbs, echoing findings by Couture et al. (2024), who document similar dynamics when the income of top earners rises.

More broadly, this paper contributes to debates on sustainability and environmental equity in urban growth and policy design. It examines how shifts in worker and busi-

⁴This complements Chinese studies such as Fan et al. (2023), Hua et al. (2023), and Feng et al. (2025).

⁵See also Dharmasankar & Yoo (2023), who document business relocation from Seattle to nearby areas.

ness locations affect emissions and air quality, explicitly quantifying impacts through the traffic channel. This analysis builds on studies such as Glaeser & Kahn (2010), Gaigné et al. (2012), and Carozzi & Roth (2023), who examine population density, pollution, and greenhouse gas emissions in broader contexts. It also highlights the role of urban green amenities (e.g., tree cover) in mitigating pollution exposure (Brack 2002; Nowak et al. 2014; Kahn & Walsh 2015; Xing et al. 2023). Finally, the findings demonstrate that *unintended* policy externalities can be disproportionate and amplify environmental inequality, adding to studies such as Shapiro & Walker (2021), Hernandez-Cortes et al. (2023), and Currie et al. (2023), which document the *intended* impacts of direct environmental regulations.

2.2 Background

This section first discusses the history and core motivations behind minimum wage policies—namely, improving the welfare and living standards of low-wage workers. To outline the mechanisms through which place-based policies, such as local minimum wage ordinances, can influence traffic, air quality, and worker well-being by redistributing workers across space, I then review empirical evidence on the broader impacts of minimum wage policies, highlighting employment, price, income, and relocation effects. Lastly, I present a conceptual framework to illustrate why the spatial effects of these policies on workers are theoretically ambiguous. Neighborhood-specific amenities, such as access to business establishments and commuting options, are important determinants of the residential distribution of workers.

2.2.1 Minimum Wage Policy

The U.S. minimum wage (MW) policy aims to enhance the welfare of vulnerable workers by protecting them from poverty and ensuring a minimum standard of living. In 1913, the Women’s Trade Union League in Massachusetts facilitated the passage of the first state-level MW law in U.S. history. In 1938, the Fair Labor Standards Act (FLSA) established the first federal minimum wage at 25 cents per hour—approximately \$4.50 in 2023 dollars. Since then, many states have enacted their own MW laws, and the FLSA has undergone numerous amendments. The federal minimum wage was last raised to \$7.25 on July 24, 2009, and has remained unchanged since. As of April 2025, 30 states have adopted MW levels exceeding the federal minimum.

City-level minimum wages were introduced to provide additional protection for workers in large urban areas, where labor costs are higher and federal or state MWs are often non-binding. The District of Columbia was the first to implement a local MW, set at \$5.25 in 1993, followed by many other cities. During the 2010s, the number of cities with their own MW laws grew to over 50, including New York City, Chicago, and Seattle.

Changes in the minimum wage can lead to the relocation of people and economic activity. These effects arise because minimum wage policies influence labor costs, price levels, workers’ income, and local employment prospects. On the business side, higher labor costs may incentivize firms—particularly in low-wage industries—to relocate to adjacent, lower-cost areas (Dharmasankar & Yoo 2023). For workers, the effects are twofold. First, while minimum wage increases often raise earnings for low-wage workers (Dube et al. 2007; Dube et al. 2010; Dube et al. 2016; Jardim et al. 2022), they may also affect job availability, hiring, turnover, and employment flows (Katz & Krueger 1992; Card & Krueger 1994; Dube et al. 2007; Dube et al. 2010; Dube et al. 2016; Cengiz et al. 2019; Dube & Lindner 2021; Jardim et al. 2022). Second, wage gains are often

accompanied by higher local prices, as a substantial portion of wage hikes is passed on to consumers in the form of higher costs for goods, services, and rents (Dube et al. 2007; MaCurdy 2015; Allegretto & Reich 2018; Leung 2021; Borg et al. 2022). Depending on the relative magnitudes of these channels, high-MW areas may become more or less attractive to different groups of workers, resulting in either residential sorting of workers toward more affordable neighborhoods or workers and firms switching workplaces to regions with higher or lower wages (Cadena 2014; McKinnish 2017; Minton & Wheaton 2022; Dustmann et al. 2022). Together, these impacts may reshape the spatial distribution of people and economic activity and alter traffic conditions.

2.2.2 Conceptual Framework for Residential Sorting

To conceptualize the channels through which place-based policies affect residential choices, I construct a standard heterogeneous-worker model and perform comparative statics on the residential equilibrium, following Gaigné et al. (2012). The key takeaway is that the effect of a wage change on workers' residential locations is ambiguous when rents are endogenous, since the outcome depends on how amenities and commuting times vary across neighborhoods.

The City

Consider a generic city with a set of locations N . These locations form a connected *network*, allowing households to travel between any two *points* $x_1, x_2 \in N$ with distance $d(x_1, x_2)$, where $d(\cdot)$ is a metric defined on N .⁶

Each worker makes two spatial choices: a residential location $x \in N$ and a workplace location $y \in N$. By residing at x , the worker derives amenity utility $b(x)$ and pays rent

⁶The metric need not be Euclidean. Common choices include Manhattan or Minkowski distances, which are widely used in residential sorting models.

$r(x)$ for each unit of land consumed. By working at y , the worker incurs commuting time, which constitutes a commuting cost in this model—that is, the opportunity cost of time. This assumption is consistent with empirical evidence that lower-income workers tend to choose less expensive but slower commuting options (*Bureau of Transportation Statistics, 2012*). Let $t(x, y) \in [0, 1)$ denote the share of time spent commuting for a worker who lives at x and works at y , where each worker is endowed with one unit of time per period. The worker’s effective working time is therefore $(1 - t(x, y))$. I assume that $t(x, y)$ is increasing in $d(x, y)$.

Workers’ Residential Decisions

Since the focus is on residential sorting, I simplify by assuming that the workplace location y is exogenous.⁷ Workers are endowed with a heterogeneous wage ω . Each worker faces a two-step decision problem: an outer choice of residential location $x \in N$, and an inner choice of consumption bundle. The inner problem is as follows. A worker residing at x and working at y with wage ω has an effective income of $\omega[1 - t(x, y)]$, accounting for time lost to commuting. The worker allocates this income between two goods: a numeraire good q and land consumption h . The numeraire is uniformly priced across locations, while the price of land $r(x)$ varies by location. Note that $r(x)$ effectively reflects the price of a consumption bundle, including local prices for groceries, restaurants, and other services. The budget constraint is therefore $q + r(x)h = \omega[1 - t(x, y)]$.

Given residence x and endowments (ω, y) , the worker chooses q and h to maximize utility:

$$\max_{q, h} U(b(x), q, h) = b(x) (q^{1-\mu} (h - \bar{h})^\mu) \quad (2.1)$$

⁷At micro-geographic scales, workplace locations are establishment-determined, and workers cannot precisely choose where to work (Wasmer & Zenou 2006). There is also extensive evidence of business relocation in response to minimum wage changes. Accounting for workplace changes adds complexity without yielding additional insights here.

$$\text{s.t. } q + r(x)h = \omega[1 - t(x, y)]$$

The utility function represents Stone-Geary preferences, as in Fretz et al. (2022), Gagné et al. (2022), and Gaubert & Robert-Nicoud (2025).⁸ Solving for utility maximization yields:

$$\begin{cases} q^*(x, \omega, y) = (1 - \mu) \cdot [\omega(1 - t(x, y)) - r(x)\bar{h}] \\ h^*(x, \omega, y) = \bar{h} + \mu \cdot \frac{\omega[1 - t(x, y)]}{r(x)} \end{cases} \quad (2.2)$$

Substituting back, the indirect utility is:

$$V(x, \omega, y) = (1 - \mu)^{1-\mu} \cdot \mu^\mu \cdot b(x) \cdot \omega[1 - t(x, y)] \cdot r(x)^{-\mu} \quad (2.3)$$

Equilibrium Rent

Because workers optimize h and q conditional on their residence, the core problem reduces to choosing a location x . The relevant equilibrium is a residential equilibrium, where the rent function $r(\cdot)$ adjusts so that no worker has an incentive to relocate. I focus on how a worker's wage ω affects their residential choice. As noted in Section 2.2.2, workplace location y is treated as exogenous. At equilibrium, a worker with wage ω solves:

$$x^*(\omega) = \arg \max_{x \in N} V(x, \omega) \quad (2.4)$$

where $V(x, \omega) = (1 - \mu)^{1-\mu} \cdot \mu^\mu \cdot b(x) \cdot \omega[1 - t(x)] \cdot r(x)^{-\mu}$ suppressing y for simplicity.

The dual problem is:

$$\omega^*(x) = \arg \max_{\omega} \Phi(\omega, x, U) \quad (2.5)$$

⁸Stone-Geary generalizes Cobb-Douglas by incorporating minimum (subsistence) consumption levels— $\bar{h}$ in this case. Cobb-Douglas implies constant elasticity of housing demand, which contradicts empirical evidence (Albouy 2016).

where $U = V(x^*(\omega), \omega)$ and

$$\Phi(\omega, x, U) = \left(\frac{(1 - \mu)^{1-\mu} \cdot \mu^\mu \cdot b(x) \cdot \omega [1 - t(x, y)]}{U} \right)^{1/\mu} \quad (2.6)$$

Here, $\Phi(\omega, x, U)$ is the bid-rent function—the maximum rent a worker with wage ω is willing to pay for x while attaining utility level U . By duality, the equilibrium rent satisfies $r^*(x^*(\omega)) = \Phi(\omega, x^*(\omega))$.

Residential Sorting in Response to Wage Changes

With endogenously determined rents, the total effect of wage changes on residential choice x follows from the implicit function theorem applied to:

$$\frac{\partial \Phi(\omega, x, U)}{\partial \omega} = 0 \quad (2.7)$$

which yields:

$$\frac{dx^*}{d\omega} = - \frac{\frac{\partial^2 \Phi(\omega, x, U)}{\partial \omega^2}}{\frac{\partial^2 \Phi(\omega, x, U)}{\partial \omega \partial x}} \quad (2.8)$$

The numerator is negative by the second-order condition, but the sign of the denominator is ambiguous. Thus, the direction of wage effects on residential location is theoretically ambiguous.

Proposition 1 (Ambiguous Residential Effect of Wage Changes) *At residential equilibrium, the sign of $\frac{dx^*}{d\omega}$ is ambiguous. Specifically,*

$$\frac{dx^*}{d\omega} > 0 \quad \text{if and only if} \quad \frac{\bar{h}}{(h^*)^2} \left[\frac{t(x, y)}{b(x)} \cdot \frac{\partial b(x)}{\partial x} + (1 - \mu) \cdot \frac{\partial t(x, y)}{\partial x} \right] > 0 \quad (2.9)$$

When both the amenity function $b(x)$ and commuting time $t(x)$ are linear in x , the effect of ω on $x^*(\omega)$ is unambiguous. In practice, however, amenities and commuting

times vary non-linearly across space. For example, amenities may be high both near urban centers and in select suburban neighborhoods (e.g., those with strong business access), and commuting times may vary due to uneven road networks and transit options. In such cases, the sign of Expression 2.9 varies across space. Higher wages may lead some workers to move closer to work, while others may move farther to access better amenities. Thus, the residential effects of wage changes depend critically on the spatial distributions of amenities and commuting costs. Empirical analysis is therefore essential to identify the realized impacts. In particular, the model suggests that policy effects are likely to vary with amenities and commuting time, which I examine in Section 2.6.3.

2.3 Data and Summary Statistics

In this section, I first introduce the main variables and describe how they are constructed from various data sources. I then present summary statistics for all variables.

2.3.1 Traffic

Traffic outcomes describe the traffic conditions *surrounding* each locality, encompassing total volume, speed, and total travel time on major highways. All variables are aggregated to annual hourly averages at the block-group level for the years 2011 to 2019, based on the California Department of Transportation’s Performance Measurement System (PeMS) from Caltrans. The use of this data source restricts the sample to Californian cities only.

Vehicle miles traveled refers to the average total miles driven by all cars within a fixed time period, calculated at the station (road segment) level. A traffic station refers to a traffic sensor designated to monitor one highway segment. All cars passing through the sensor are assumed to travel the entire segment (the median segment length in this

sample is 0.494 miles). This variable is defined as the number of cars passing through a road segment during a given time period, multiplied by the segment’s length in miles, as in Knittel et al. (2016). Average speed measures the average speed of traffic flow on each road segment. *Vehicle hours spent* is defined as the total miles traveled by all cars divided by the average speed (i.e., *vehicle miles traveled* divided by average speed), representing the total amount of time drivers spend on the road. The share of time with average speed below 60 or 50 mph approximates the likelihood of congestion and more severe congestion, respectively.

To obtain traffic outcomes at the annual block-group level, I first compute the annual average for each traffic measure, and then I use inverse-distance weighting to aggregate data from the station (road-segment) level to the block-group level.⁹ Following the recommendations of Knittel et al. (2016), I include all highway segments within 15 miles of a block group’s centroid. These variables thus reflect not only traffic conditions within a block group but also those in nearby areas. Since PeMS vehicle detector stations are predominantly located along Interstate Freeways (e.g., I-5, I-80), U.S. Routes (e.g., 101), and State Routes (e.g., CA-57), these traffic measures primarily reflect changes in freeway traffic. Traffic on local roads and bypasses is not captured.

In terms of data quality, missing values due to non-operation of traffic sensors are expected to have a limited impact on the results. Caltrans applies a sequence of linear regression imputations based on neighboring stations, using local coefficients (estimated from nearby stations), global coefficients (from all stations), and temporal medians to replace missing values. The system also reports the share of observed data as a quality measure. As discussed in Knittel et al. (2016), the local imputation method performs well in approximating true traffic trends. When the imputed data are approximately

⁹For example, if there are only three road segments near block group A, each at 1, 5, and 7 miles away, the traffic condition in A is the average of the three, with weights 1, 1/5, and 1/7, respectively.

unbiased, using such data yields conservative estimates of the true effects, given that approximation introduces noise that can be treated as classical measurement error, which typically attenuates estimated coefficients toward zero. This is likely the case given the use of time- and location-specific fixed effects, which absorb part of the errors that are constant at the same level as the fixed effects. I conduct a robustness check using only unimputed data and find similar results with slightly larger magnitudes. They are close to the main results and are therefore omitted.

2.3.2 Air Quality

Air quality is measured as the annual average levels of $PM_{2.5}$ and NO_2 at the block-group level, derived from Van Donkelaar et al. (2021) and Cooper et al. (2020). $PM_{2.5}$ and NO_2 are critical pollutants for human health, and ground transportation is a major source of both. Particulate matter is partially generated by fuel combustion, friction during acceleration and braking, and atmospheric reactions involving traffic-related emissions (Li & Managi 2021). Exposure to fine particulate matter has been linked to respiratory issues (Xing et al. 2016, Fan et al. 2016), chronic diseases (Wang et al. 2015, Leonetti et al. 2024), and mortality (Burnett et al. 2018, Woodruff et al. 2006, McDuffie et al. 2021). NO_2 primarily originates from fuel combustion (EPA, 2024). Zhong et al. (2017) find that changes in traffic and congestion significantly influence NO_2 concentrations and are associated with increases in emergency medical cases.

The dataset from Van Donkelaar et al. (2021) provides annual $PM_{2.5}$ concentrations at a fine spatial resolution (0.01° , approximately 1 km). The estimates integrate satellite-based aerosol optical depth (AOD) data with chemical transport models and ground observations, enabling comprehensive global assessments of ground-level $PM_{2.5}$, including in areas lacking monitoring infrastructure. This dataset has been widely used

in studies examining the relationship between particulate matter and health (Xue et al. 2022, Rentschler & Leonova 2023, Han et al. 2023). The dataset from Cooper et al. (2020) provides annual ground-level NO_2 concentrations for 2005–2019 by integrating high-resolution data from the TROPOMI satellite with long-term records from the OMI instrument. Following Liddell et al. (2025), I convert grid-level pollutant concentrations to block-group-level values by overlaying census block-group boundaries on air quality data grids.


2.3.3 Spatial Distribution of Workers and Commute Patterns

As discussed in Section 2.2, minimum wage ordinances may affect workers’ spatial distribution on both the residential and workplace sides. I use Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics to measure the spatial distribution of workers and their commute patterns. I explicitly differentiate between workplace and residence, as the policy affects each through different channels. I count the total number of workers by workplace block group (*employment size*) to capture direct policy impacts through business and employment channels. Next, I report the total number of workers by residential block group (*resident size*) to assess impacts through income and residential sorting channels. I also construct a *employment-resident ratio*, defined as the difference between the logarithm of employment size and the logarithm of resident size, to quantify the relative attractiveness of locations as workplaces versus residences. A higher *workers-residents ratio* implies that a location is more attractive as a workplace relative to a residence.

To jointly capture workplace and residence effects, I report commuting patterns, including commuting time, distance, and types, based on the pairing of workplace and residence locations. Specifically, I classify workers into four commuting types, depending

on whether their residence and workplace fall within or outside the policy boundary. I then count the number of workers in each category. These categories collectively sum to the total number of workers by workplace or residence. I use these disaggregated patterns to quantify changes in commuting arrangements in response to the policy. Figure 2.1 illustrates the construction of the employment size, resident size, and commuting type variables and provides a conceptual framework for the paper.

Figure 2.1: Conceptual Framework and Commuting Variables Construction



Notes: This diagram illustrates how a minimum wage may affect the spatial distribution of workers, their commuting patterns, as well as traffic patterns and air quality. The three text boxes in the middle also demonstrate the relationship between location and commuting variables. The four types of commuting arrangements are reported in the middle green text box.

For commuting distance, I calculate the median geographic distance for all workers employed in the same workplace block group. Distances are computed using the centroids of residential and workplace block groups via the Haversine formula. To approximate

median driving distances and times, I match block groups to census tracts and link them to cross-sectional Census Tract Origin-Destination Driving Time (2017) data from the Urban Institute. Routing is based on the Open Street Map Routing Machine (OSRM) using Open Street Map North America data. Because these data are cross-sectional, they do not reflect changes in commuting time or distance due to evolving road infrastructure or traffic conditions over time.

Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics is an administrative dataset that provides annual job/worker counts by workplace census block and by residence census block, as well as by workplace–residence census block pair. I use the primary-job file to focus on main employment and restrict the analysis to the pre-COVID period, when remote work was less prevalent. The sample includes only workers whose residence and workplace fall within the same commuting zone. For heterogeneity analysis, I further use information on industry and worker demographics to categorize workers by Hispanic origin, education level, and industry to identify groups most affected by the policy.

2.3.4 Policy Variable

The main policy variable is the city-level minimum wage for each city-year. Prior to the adoption of city-level ordinances, minimum wage levels were set by federal and state regulations. The Department of Labor documents these policy changes since 1938. The Inventory of US City Minimum Wage Ordinances by the UC Berkeley Labor Center provides city-level minimum wage data, including adoption dates, initial wage levels, and subsequent updates.

I use the level of the minimum wage, rather than an indicator for policy adoption, to capture richer variation in policy intensity. Unlike federal or other state policies, which


change infrequently, California and most cities in my sample adjust minimum wage levels periodically. As a result, minimum wages vary both over time and across cities. By including time fixed effects to absorb state-level adjustments and other macroeconomic trends, the level-based measure compares policy cities with relatively higher versus lower minimum wage levels. Table 2.1 lists all policy cities and their respective ordinances. Because of data availability, only Californian cities are included in this paper. Figure 2.2 shows trends in average minimum wage levels over time.

Table 2.1: Summary statistics on CMWs and adjustments

City	Implement date	Avg. policy adjustment duration (in days)	Avg. increase	Max. increase	Min. increase
Richmond	2015-01-01	365	1.10	1.92	0.60
San Diego	2015-01-01	487	0.75	1.00	0.50
Sunnyvale	2015-01-01	365	1.50	2.00	0.70
Oakland	2015-03-01	350	0.96	3.25	0.30
Mountain View	2016-01-01	365	1.34	2.00	0.65
Palo Alto	2016-01-01	365	1.25	1.50	1.00
Santa Clara	2016-01-01	365	1.25	2.00	0.10
El Cerrito	2016-07-01	305	1.25	1.60	0.65
Los Angeles	2016-07-01	365	1.06	1.50	0.50
Malibu	2016-07-01	365	1.06	1.50	0.50
Pasadena	2016-07-01	365	1.06	1.50	0.50
Santa Monica	2016-07-01	365	1.06	1.50	0.50
Berkeley	2016-10-01	334	1.15	1.53	0.59
Cupertino	2017-01-01	365	1.67	2.00	1.50
San Mateo	2017-01-01	365	1.67	2.00	1.50
San Jose	2013-03-01	274	1.50	1.50	1.50
San Leandro	2017-07-01	365	1.17	1.50	1.00
Belmont	2018-07-01	184	1.25	1.50	1.00

Notes: There are 17 California cities in total. The average policy adjustment duration shows, for each city, how many days a given minimum wage level remains in effect on average. As previously noted, minimum wages are typically implemented via ordinance and are adjusted approximately annually. I report increments only, as the minimum wage increases monotonically over time.

Figure 2.2: Time Trends of Average Minimum Wage Levels with and without City-Level Ordinances



Notes: This figure reports the time trends of average minimum wage levels in the sample period across all policy cities, with and without considering city-level ordinances. The level of minimum wage without city ordinances is later used in the simulation in Section 2.7.1 to generate counterfactuals.

2.3.5 Sample and Summary Statistics

The policy sample includes 17 California cities that implemented city-level minimum wage increases between 2010 and 2018. I only include California cities due to data restrictions. I restrict the sample to cities with such ordinances, as the remaining cities might not be comparable (Dube & Lindner 2021). I include census block groups located within the policy boundaries to capture direct policy impacts. I separately consider block groups that are outside the policy limits but within the same commuting zone and no more than 50 miles from the policy city. This sample is used to capture policy spillovers and general equilibrium effects.

Table 2.2: Summary Statistics

	Mean	S.D.	Min.	25th	75th	Max.
Policy and Location Variables						
Minimum Wage Level	10.17	2.21	8.00	8.00	12.00	15.65
Located Inside Policy Limits	0.35	0.48	0.00	0.00	1.00	1.00
Air Quality						
$PM_{2.5}$	10.84	2.31	3.99	9.01	12.57	17.97
NO_2	0.36	0.14	0.00	0.27	0.42	1.30
Traffic						
Vehicle Mile Traveled	1715.50	1269.97	523.91	1224.84	1,833.51	14111.51
Vehicle Hour Spent	30.32	21.75	8.46	21.50	34.29	262.29
Average Speed (mph)	62.04	2.45	46.28	60.66	63.99	71.67
Probability of Congestion	0.20	0.10	0.00	0.13	0.25	0.98
Probability of Severe Congestion	0.09	0.06	0.00	0.05	0.12	0.83
Location and Commute						
Employment Size	640.72	2568.25	1.00	65.00	431.00	142320
Resident Size	636.54	399.07	1.00	371.00	804.00	9477
Employment-Resident Ratio	3.33	1.45	-1.88	2.34	4.22	10.11
N Workers From Policy Cities	205.69	1,039.87	0.00	10.00	117.00	44517
N Workers From Neighboring Areas	302.29	1,322.05	0.00	17.00	197.00	62526
Median Geographical Miles	9.45	5.73	0.17	6.44	11.41	397.37
Median Driving Miles	12.32	6.12	0.00	8.60	14.80	125.30
Median Driving Distance > 5 Miles	0.73	0.13	0.00	0.65	0.82	1.00
Median Driving Distance > 10 Miles	0.54	0.14	0.00	0.44	0.64	1.00
Median Driving Distance > 15 Miles	0.40	0.14	0.00	0.30	0.49	1.00
Median Driving Minutes	20.05	8.38	0.00	15.00	23.40	158.10

The unit of observation is a census block group by calendar year. The sample includes block groups located within 50 miles of the policy boundary and within the same commuting zone. Vehicle miles traveled and vehicle hours spent refer to the total miles (or hours) traveled by cars through highway segments per hour. For each census block group, I use inverse distance weighting to aggregate data from all highway segments within 15 miles. The variable “Block Groups Inside the Policy Limits” is a binary indicator equal to 1 if block group i lies within the policy limits. I divide the sample based on this indicator. The total number of observations is 123,294. Employment size does not match resident size because the sample only includes BGs within 50 miles, and some workers may work or reside elsewhere.

For block groups that are outside the policy limits but nearby, I assign them the minimum wage of the nearest policy city and measure how their outcome variables respond to the ordinances. To differentiate main policy impacts from spillover and general equilibrium effects, I estimate the policy impacts separately for block groups located within versus outside policy limits. I prioritize distance, as Dharmasankar & Yoo (2023) and Jardim et al. (2024) find that distance plays a pivotal role in spillover effects. To address the concern that one location might be affected by several policy cities, I construct alternative policy variables, defined as the average minimum wage level among nearby policy cities (see Appendix Table B.4), which yield results similar to the main specification.

The sample spans 2011 to 2019, focusing on the five years preceding and the three years following initial city-level minimum wage implementations. The analysis is conducted at the census block group by calendar year level, yielding a total of 123,294 observations (42,822 inside policy limits and 80,486 outside). Summary statistics are reported in Table 2.2.

2.4 Empirical Strategy

I begin by outlining an idealized experiment to illustrate the logic of the analysis before turning to the econometric details. An omnipotent policymaker would randomly assign minimum wage levels to each city and estimate the policy impact using ordinary least squares. The empirical strategy aims to approximate this experiment but faces two main challenges. First, the timing and levels of minimum wages are not randomly assigned; they may respond to local economic and political conditions. Second, block groups may differ in ways that are fixed over time but nonetheless affect traffic, air quality, or commuting patterns, such as geographic layout or historical infrastructure investment.

To identify the causal impact of city-level minimum wage increases on workers’ spatial distribution and commuting patterns, as well as the unintended policy impacts on traffic and air quality, I follow Allegretto et al. (2011), Cadena (2014), Dube (2019), and Leung (2021) by exploiting variation in minimum wage levels across space and over time. The central idea is to compare more-treated locations to less-treated ones, rather than to untreated locations.

$$y_{ict} = \beta MW_{it} + \eta_i + \eta_{ct} + \epsilon_{ict} \quad (2.10)$$

Here, i indexes census block groups, c indexes counties, and t indexes years. The policy variable MW_{it} is the statutory minimum wage level of the policy city to which block group i is assigned. Details of this assignment are described in Section 2.3.3, with robustness checks provided in the Appendix. I include block-group fixed effects (η_i) to control for all time-invariant differences across block groups (e.g., geographic conditions and infrastructure) and county–year fixed effects (η_{ct}) to absorb shocks or trends common to all locations in the same county in a given year (e.g., countywide macroeconomic conditions, labor market policies, or transportation and pollution regulations).

By the Frisch–Waugh–Lovell theorem, this setup residualizes MW_{it} with respect to η_i and η_{ct} . The identifying variation comes from deviations in the minimum wage from the county–year average, net of each block group’s long-run mean. That is, the comparison is between block groups in the same county and year that differ in their exposure to minimum wage levels, and between the same block groups over time as their relative minimum wage exposure changes. The counterfactual is block groups in the *same county and year* but exposed to a lower minimum wage, holding constant time-invariant differences and county–year shocks. Thus, β captures how outcomes change when the minimum wage rises by one dollar *relative to other policy cities in the same county at the same time*, under the identifying assumption that, conditional on the fixed effects—these

city-specific deviations in minimum wage are uncorrelated with unobserved determinants of the outcomes.

I conduct several robustness checks and placebo tests to address identification concerns, summarized briefly here and detailed in Section 2.5.3. First, I account for the endogeneity of policy implementation and adjustment, as well as dynamic effects, by adding leads and lags. Appendix Table B.2 shows that allowing for dynamic treatment does not alter the main estimates. Second, to test the validity of the specification, I replace county-year fixed effects with county-subdivision-year fixed effects and county-specific linear trends, finding consistent results. I also implement the heterogeneity-robust difference-in-differences estimator of de Chaisemartin et al. (2022), which accommodates continuous treatments, heterogeneous effects, and nonlinearities by comparing switchers to stayers under conditional parallel trends. Placebo tests also provide support for the parallel pre-trends assumption (Appendix Table B.3). Third, to address possible Stable Unit Treatment Value Assumption violations, I redefine treatment as the weighted average of all nearby policy cities rather than only the closest one, with results reported in Appendix Table B.4. Fourth, following Cengiz et al. (2019) and Dube & Lindner (2024), I estimate a stacked event-study design that forms control groups from later-treated cities. This approach is robust to heterogeneous effects and staggered adoption (Dube et al. 2023), with results reported in Appendix Tables B.5–B.6. Section 2.5.3 provides a detailed discussion of these robustness checks.

To account for population dynamics in conjunction with changes in congestion and air quality, traffic and air quality regressions are weighted using work-adjusted population estimates. Specifically, I use the 2010–2014 5-year American Community Survey (ACS) as a benchmark to represent population weights for 2012, and then adjust these weights annually based on changes in *resident size* (number of workers by residence block

group) from Longitudinal Employer-Household Dynamics (LEHD) data.¹⁰ This approach captures the relocation of workers but not other household members, and thus underestimates total population movement. To assess its implications, I compare results using static (ACS) versus worker-adjusted (ACS+LEHD) weights. Both sets of estimates have the same sign, while those with worker-adjusted weights are larger in magnitude, suggesting that population movement is concentrated toward more exposed areas. Thus, when accounting for additional movements of household members, the effects are plausibly even larger, making the current approach conservative.

2.5 Results

In this section, I present evidence that place-based minimum wage ordinances worsen traffic congestion in both policy cities and neighboring areas, with the latter experiencing increased vehicle miles and hours traveled, as well as worsened air quality, despite not being directly targeted by the policy. I then explore the underlying mechanism, showing that these ordinances alter the spatial distribution of workers, pushing them toward neighboring areas, particularly in terms of residential location. These shifts have direct consequences for commuting, with more workers traveling from and to surrounding areas.

2.5.1 Traffic and Air Quality

Table 2.3 reports the policy impacts on traffic and air pollutant concentrations separately for census block groups inside (column 1) and outside (column 2) policy limits. Note that instead of clustered standard errors, I report wild cluster bootstrap **P-values** in square brackets. For policy cities, traffic impacts are mixed. Vehicle miles traveled and

¹⁰For example, if the ACS benchmark gives 1,000 residents for block group A in 2012, and LEHD reports 20 workers moving out from 2012 to 2013 and 50 moving in from 2012 to 2014, then the population weights for this block group are 980 in 2013 and 1,050 in 2014.

vehicle hours spent show small, insignificant increases, while average traffic speed falls by 0.1819 miles per hour and the probability of more severe congestion (average speed < 50 mph) rises significantly by 0.42 percentage points. In terms of air quality, there is limited evidence that the policy affects ambient concentrations within policy cities.

Table 2.3: Policy Impacts on Traffic and Air Quality


	Inside Policy Limits	Outside Policy Limits
<i>Panel A. Traffic</i>		
Vehicle Miles Traveled	0.0034 [0.3206]	0.0330** [0.0080]
Vehicle Hours Spent	0.0077 [0.2485]	0.0384* [0.0140]
Average Speed	-0.1819* [0.0240]	-0.1567* [0.0180]
% Congestion	0.0085 [0.1122]	0.0073** [0.0020]
% More Severe Congestion	0.0042* [0.0120]	0.0039* [0.0120]
<i>Panel B. Air quality</i>		
$PM_{2.5}$	0.0013 [0.1904]	0.0074* [0.0341]
NO_2	-0.0001 [0.4509]	0.0061* [0.0100]
N	42698	80189

Notes: Estimations are based on equation 2.10, separated for census block groups within (column 1) and outside (column 2) policy boundaries. $PM_{2.5}$, NO_2 , vehicle miles, and vehicle hours are in logarithmic forms, while the rest of the outcomes in this table are in levels. % Congestion (% More Severe Congestion) represents the proportion of time during which the average speed falls below 60 (50) mph. Controlling for block group and county-by-year fixed effects. **Wild clustered bootstrap P-values** are reported in square brackets. *, **, and *** indicate significance at the 0.05, 0.01, and 0.001 levels, respectively.

In contrast, minimum wage ordinances significantly increase traffic usage and worsen congestion and air quality in block groups outside policy limits. A one-dollar higher

minimum wage leads to substantial increases in traffic volume (+3.30%), travel time (+3.84%), lower average speeds (−0.1567 mph), and significantly higher congestion probabilities (+0.39 to +0.73 percentage points).¹¹ As a result, pollutant concentrations rise by 0.74% for $PM_{2.5}$ and 0.61% for NO_2 . To visualize these impacts, Figure 2.3 displays changes in $PM_{2.5}$ concentrations before and after policy implementation, comparing areas within and outside policy limits. The inner circle (policy cities) shows null to small reductions in $PM_{2.5}$, while neighboring areas experience increases. Similar figures for Northern California and for NO_2 are reported in Appendix Figure B.1.

Figure 2.3: Map of $PM_{2.5}$ Concentration Changes in Southern California



Notes: This figure illustrates the average change in $PM_{2.5}$ concentrations in Southern California before and after the city-level implementation of the minimum wage policy. The analysis spans up to five years before and three years after the policy implementation. The data are presented in logarithmic form, de-trended by calendar year, and calculated at the census block group level. The inner (purple) lines delineate the city limits subject to the minimum wage policy, while the outer (black) lines define the boundaries of the analysis sample. The analysis sample includes all census block groups within the same commuting zone and located no more than 50 miles from the policy cities. Areas situated between the two sets of lines are designated as “Outside policy limits” in the analysis. Appendix Figure B.1 reports results for other pollutants and study areas.

¹¹Appendix Table B.7 shows no significant changes in public transit ridership.

2.5.2 Spatial Distribution of Workers and Commute Patterns

To explain the traffic and air quality impacts from the previous subsection, Table 2.4 presents the impacts of minimum wage ordinances on workers' spatial distribution by employment and residential location. The left column reports policy impacts on the number of workers who work (employment size) or reside (resident size) in policy cities, while the second column reports impacts for those who live or work outside policy cities. For policy cities (column 1), both employment (-0.59%) and resident size (-0.97%) decline, with the latter falling more sharply, suggesting that fewer workers both work and, more significantly, reside in policy cities. Conversely, neighboring areas outside policy limits experience modest increases in employment ($+0.14\%$) and resident size ($+0.37\%$). Together, these results suggest shifts in the working population from policy cities to neighboring areas, particularly in terms of residential location. The employment-resident ratio also rises in policy cities ($+0.0101$) but falls in neighboring areas (-0.0153), indicating that policy cities become relatively more attractive as places of work, whereas surrounding areas become relatively more attractive as places of residence. In Section 2.6.2, I further examine employment impacts by worker characteristics and situate my estimates in the broader literature.

These shifts in residential preferences are corroborated by rent estimates. For this exercise, I use the Zillow Observed Rent Index (All Homes Plus Multifamily Time Series) and estimate how average rents respond to minimum wage levels using the same specification as Equation 2.10. The results show significant rent increases, nearly twice as large in policy cities ($+0.164\%$) as in neighboring areas ($+0.085\%$ per \$1 MW). This reflects the pass-through of higher labor costs to consumers, as documented by MaCurdy (2015) and Borg et al. (2022). Rising living costs push workers to seek more affordable housing in neighboring areas.

Table 2.4: Policy Impacts on the Spatial Distribution of Workers and Rents

	Inside Policy Limits	Outside Policy Limits
Employment size	-0.0059* [0.0120]	0.0014 [0.2064]
Resident size	-0.0097** [0.0040]	0.0037 [0.5411]
Emp.-Res. Ratio	0.0101 [0.1122]	-0.0153 [0.2104]
Rents	0.00164*** [0.0010]	0.00085* [0.0350]
N	42808	80486

Notes: Estimations follow Equation 2.10, estimated separately for census block groups within (column 1) and outside (column 2) policy boundaries. **Wild clustered bootstrap P-values** are reported in square brackets. Employment and resident sizes are in logarithmic form. The rent index measures observed rents within the 35th–65th percentile of the market at the zip-code-month level and enters regressions in logarithmic form. The number of observations is 4,649 and 6,719 for the rent estimates. Following Borg et al. (2022), I estimate the effect of minimum wage on logarithmic rents using a first-difference model to increase statistical power, as rents are serially auto-correlated. All specifications include block group and county-by-year fixed effects (county-by-month for rents).

The observed divergence in impacts on employment and residence implies changes in commuting patterns, as commuting connects residential and workplace locations. Table 2.5 reports policy impacts on commuting types and distances. The outcome variables in Panel A measure the number of workers in four commute categories, described in Section 2.3.3 and illustrated in Figure 2.1. The top-left cell describes workers whose workplace and residence are both inside the policy limits; the top-right cell describes workers who live inside but work outside; the bottom two cells describe workers who reside outside the policy limits (and work either inside or outside). Panel B summarizes impacts on commute lengths. The construction of these outcomes is detailed in Section 2.3.3 and illustrated in Figure 2.1.

Table 2.5: Policy Impacts on Commute Prevalence and Lengths, by Location of Work

	Workplace Inside Policy Limits	Workplace Outside Policy Limits
<i>Panel A. Commuting Arrangement: N Workers residing in</i>		
Policy Cities	-0.0072* [0.0261]	-0.0026* [0.0100]
Neighboring Areas	-0.0036 [0.2044]	0.0023 [0.2184]
<i>Panel B. Commuting Distance and Time</i>		
Median Geographical Distance	0.0498* [0.0477]	-0.2773** [0.0050]
Median Driving Distance (MDD)	0.0691* [0.0477]	-0.2510*** [0.0000]
MDD > 5 Miles	0.0022* [0.0276]	-0.0002 [0.2814]
MDD > 10 Miles	0.0017* [0.0452]	-0.0009 [0.1156]
MDD > 15 Miles	0.0011 [0.1005]	-0.0026** [0.0050]
Median Driving Minutes	0.0736* [0.0251]	-0.3421*** [0.0000]
N	42808	80486

Notes: Estimations follow Equation 2.10, with results reported separately for block groups within (column 1) and outside (column 2) the policy boundaries. Since outcomes in Panel A are constructed based on workplace block groups, Column 1 reflects impacts for workers employed within policy limits, while Column 2 reflects those employed outside. Outcomes in Panel A are the number of workers by four commute types (workplace within/outside $\times$ residence within/outside policy limits) in logarithms, while those in Panel B are in levels. For example, the top left cell describes the number of workers with both workplace and residence inside policy limits, while the top right describes those who live inside but work outside. Due to the cross-sectional nature of travel time data, estimated impacts on commute duration are location-driven and do not reflect changes in traffic conditions. All specifications include block group and county-by-year fixed effects. **Wild clustered bootstrap P-values** are reported in square brackets.

The first column of Panel A shows changes in commute patterns for workers employed

in policy cities. There are significantly fewer (-0.72%) workers who both live and work in policy cities. The reduction is smaller (-0.36%) and insignificant for workers employed in policy cities but residing outside. Among all workers employed in policy cities, the median one-way commute distance increases by 0.0498 geographic miles (or 0.0691 driving miles). Commute time also rises significantly, and the share of workers commuting five miles or more increases by 0.22 percentage points.

The second column reports analogous estimates for workers employed outside policy cities. Commuting from policy cities declines significantly by 0.26%, while the number of workers both living and working outside increases, though not significantly. Fewer external commuters (residing in policy cities) and more internal commuters (residing outside) together lead to a significant reduction in commute distance and time.

2.5.3 Robustness Check

As discussed in Section 2.4, the empirical approach leverages treatment intensity and compares outcomes in block groups exposed to higher versus lower levels of city minimum wage, under the identifying assumption that—conditional on the fixed effects—these city-specific deviations in minimum wage are uncorrelated with unobservables. To corroborate the validity of this approach, I conduct the following robustness checks to address potential identification concerns.

First, a potential concern is that the aforementioned specification does not facilitate meaningful comparisons. To address this, I re-estimate the model using alternative specifications, replacing county-year fixed effects with either county-subdivision-year fixed effects or county-specific linear trends. Appendix Table B.1 reports the benchmark estimates (columns 3 and 6) along with two alternative specifications (columns 1–2 and 4–5). For air quality, the alternative estimates show more positive point estimates, but they

are not significant due to lower power. For traffic, estimates across specifications consistently indicate increases in vehicle miles traveled, vehicle hours spent, and congestion probability, and a lower average speed for block groups outside policy limits. In addition, estimates for block groups inside policy cities are roughly the same for average speed and congestion probabilities. Estimates for vehicle hours spent and vehicle miles traveled in policy cities show a reversal in sign but are not statistically significant. For commuting outcomes and the spatial distribution of workers, the results are broadly consistent across specifications.

Second, a second concern is policy anticipation, as ordinances are typically announced before implementation. To address this, I include two-year leads and lags in equation 2.10 and re-estimate the policy impacts. The results suggest similar contemporaneous effects on the spatial distribution of workers and commuting, with small increases in vehicle miles, vehicle hours, and higher ambient $PM_{2.5}$ and NO_2 concentrations in policy cities, although the estimates are insignificant. The contemporaneous impacts for block groups outside policy limits are similar, with significantly more traffic, worse air quality, shifts in the spatial distribution of workers, and reduced commuting. In terms of anticipation, most estimates are insignificant, especially for block groups outside policy limits. Yet, the employment–resident ratio and average speed indicate changes inside policy cities one year before actual implementation. For lagged effects, many estimates share the same sign as the contemporaneous impacts and are statistically significant for block groups inside policy limits.

Third, the recent literature highlights that two-way fixed-effects estimators can be biased in the presence of heterogeneous treatment effects. To address this, I follow de Chaisemartin et al. (2022) and implement the heterogeneity-robust difference-in-differences (HRDID) estimator, which accommodates continuous treatments as in my setting. This method estimates weighted slopes for *switchers* (units whose treatment

changes over time) and is robust to both heterogeneous and nonlinear effects. The comparison group comprises *stayers* (units with no change in treatment), under the conditional parallel trends assumption—that conditional on covariates, switchers and stayers would have followed similar trends in the absence of treatment. To assess this assumption, I test for placebo effects by examining whether switchers and stayers followed parallel trends prior to treatment. Appendix Table B.3 presents the main effects along with placebo estimates. Significant increases are observed in $PM_{2.5}$ and NO_2 concentrations, vehicle miles traveled, and vehicle hours for neighboring block groups. In contrast, in policy cities, HRDID shows a significant increase in NO_2 concentration and a sharp rise in the employment–resident ratio. Notably, most placebo effects are insignificant.

Fourth, violations of the Stable Unit Treatment Value Assumption (SUTVA) are a potential concern. Given that more than one policy city may be nearby, it is useful to know the impacts from all rather than only the closest one. I construct an alternative policy variable as the weighted average of minimum wages in nearby policy cities and re-estimate the policy impacts in Appendix Table B.4. Using the average minimum wage as the policy variable yields smaller effects on air quality and traffic for block groups outside policy limits, while commute changes become larger for workers employed in policy cities. SUTVA is further tested in the stacked design below, with comparison groups restricted to untreated areas.

Fifth, following Cengiz et al. (2019) and Dube & Lindner (2024), I estimate a stacked event-study design of the policy impact. Specifically, I construct comparison groups for each policy city using cities not concurrently treated but treated in later years. I also restrict the sample to cities treated no later than 2022 to ensure that later-treated cities in the main sample have appropriate controls. Each policy city and its corresponding control group constitute an *event*. I then stack the data and estimate policy effects conditional on event–year and event–block group fixed effects, ensuring that the estimates capture

only within-event variation. This method is robust to heterogeneous treatment effects and staggered policy adoption (Dube et al. 2023). Results are reported in Appendix Table B.5, using minimum wage levels as the policy variable, and in Appendix Table B.6, using a binary treatment indicator.

2.6 Heterogeneous Impacts

This section documents the uneven distribution of congestion and air quality impacts across neighborhoods, with air quality effects disproportionately concentrated in minority and low-income communities. I first show that these communities are located closer to major roadways and have less access to green amenities. As a result, they experience greater air quality impacts even under similar traffic conditions. I also show that minimum wage ordinances disproportionately affect Hispanic, low-education, and service-sector workers, but these workers do not appear to be sorting into the communities most affected by pollution. Instead, amenities and commute times appear to be more significant drivers of their residential choices.

2.6.1 Disproportionate Air Quality Impacts and the Causes

I use neighborhood characteristics from the ACS to explore heterogeneous impacts by socioeconomic status. I retrieve the share of the population with various characteristics from the 5-year ACS 2010–2014.¹² I divide the sample into four subgroups based on whether a block group ranks above or below the median in a given characteristic, and whether it lies within or outside the policy limits, as in the previous section. To facilitate


¹²Ideally, one would observe workers’ characteristics by both residence and workplace. However, due to data limitations, information is only available by residential location. Given that workers typically spend more time at home than at work, this approach approximates the main part of the air quality impacts.

comparison, I report the policy impacts for four subgroups side-by-side in the same panel. I consider six characteristics: (1) Hispanic ethnicity; (2) four-year college attainment; (3) employment in the service sector (NAICS code 72); (4) standard 9 AM–5 PM weekday schedule; (5) household income above twice the poverty line; and (6) public transit usage.

Figure 2.4 shows that increases in $PM_{2.5}$ concentrations are more prominent in neighborhoods outside the policy limits that have higher poverty rates, lower public transit access, and greater shares of Hispanic, service-sector, less-educated, or irregular-schedule workers (i.e., those who do not work a standard 9 AM to 5 PM weekday schedule). The blue circle in the top-left sub-panel represents the $PM_{2.5}$ impact for block groups within the policy limits and above the median Hispanic share. This estimate is compared to the red diamond (inside policy limits, below median), the brown triangle (outside policy limits, above median), and the green square (outside policy limits, below median). The comparison reveals substantial increases in pollutant concentrations for block groups outside the limits with high Hispanic shares.

To explain differential air quality impacts, I consider two possible explanations. First, traffic impacts may differ across neighborhoods, such that those with greater traffic also experience more severe air quality consequences. Second, certain neighborhoods may have a lower capacity to adapt to air pollution. Finally, I provide suggestive evidence that the observed impacts are not driven by emissions from other sources, such as residential heating.

Figure 2.4: Heterogeneous Impacts on Air Quality and Traffic by Neighborhood Characteristics



Notes: The top panel reports policy impacts on $PM_{2.5}$ by census block group location (inside vs. outside policy limits) $\times$ neighborhood characteristic ranking (top vs. bottom 50%). Each sub-panel corresponds to a specific characteristic, including: % Hispanic population, education, % service workers, % weekday 9 AM–5 PM workers, % poverty, and public transportation access. Characteristic rankings are based on the 2010–2014 5-year ACS. From left to right, each sub-panel reports estimates for block groups that are: (1) inside policy limits and above the median level of the characteristic; (2) inside policy limits and below the median; (3) outside policy limits and above the median; and (4) outside policy limits and below the median. The bottom panel reports the same estimates for vehicle miles traveled. Outcome variables are in logarithmic form with 95% C.I. reported.

To test the first explanation, I examine whether traffic impacts differ by neighborhood characteristics and report the results in the bottom panel.¹³ For neighborhoods outside the policy limits, impacts on vehicle miles traveled differ by block group characteristics, but these differences are not statistically significant. In Appendix Figure B.2, I further explore impacts on vehicle miles traveled and speed by time of day and day of week. The results suggest that vehicle miles traveled increase primarily during weekday mornings and evenings outside policy cities. By contrast, speed reductions within policy cities are concentrated in the morning to midday, while in neighboring areas they occur in both the early morning and late evening.

I next examine adaptation capacity by testing whether more-affected neighborhoods are closer to major roads and/or lack green space. The evidence supports this conjecture. Specifically, I measure total road length within a 3-mile radius of the block group centroid (using PeMS) and calculate the share of green space within the block group (using the National Land Cover Dataset 2012). I again divide neighborhoods into high/low groups based on each characteristic and report differences in Table 2.6. For each characteristic, I report the average for the low group, and the difference and p-value compared to the high group. Results show that vulnerable communities are surrounded by denser road networks and have significantly less vegetative cover to shield them from pollutants. Thus, even with equivalent traffic impacts, air quality impacts are expected to be larger for the low group. To directly test this, I report air quality impacts for block groups in the top versus bottom shares of nearby highways in Appendix Table B.8. Compared to the full sample, block groups in the top decile or quartile indeed experience larger impacts, while those in the bottom quartile are largely unaffected.

¹³Additional results for NO_2 , vehicle hours, average speed, and congestion probability are reported in Appendix Figures B.3–B.4.

Table 2.6: Amenity in Communities Above and Below Median Levels of Selected Characteristics

	Total Road Length			Vegetation		
	Mean	% Difference	P-Value	Mean	Difference	P-Value
<i>Panel A. Inside Policy Limits, N=5118</i>						
Hispanic	80.24	19.44	0.0183	0.059	-0.038	<0.0001
Lower Education	82.42	16.12	0.0125	0.056	-0.035	<0.0001
Service Workers	82.33	13.55	0.0023	0.060	-0.041	<0.0001
Non-Standard Work Schedule	85.93	5.30	0.0001	0.050	-0.025	0.0003
Public Transp.	84.43	6.25	0.1958	0.053	-0.024	<0.0001
Commuting Time	85.56	4.53	0.5333	0.034	0.006	0.3320
Lower Rent	80.51	17.71	0.0010	0.058	-0.036	0.0001
Poverty	80.35	19.24	0.0036	0.066	-0.048	<0.0001
<i>Panel B. Outside Policy Limits, N=8356</i>						
Hispanic	75.81	23.30	0.0007	0.149	-0.088	<0.0001
Lower Education	76.20	18.71	0.0250	0.142	-0.070	<0.0001
Service Workers	78.80	13.73	0.0204	0.136	-0.061	<0.0001
Non-Standard Work Schedule	80.01	8.78	0.0307	0.124	-0.033	<0.0000
Public Transp.	85.97	-0.80	0.5457	0.109	-0.014	0.1257
Commuting Time	87.21	-11.70	0.0002	0.077	0.059	0.0001
Lower Rent	78.91	14.53	0.0439	0.130	-0.049	<0.0001
Poverty	76.80	20.03	0.0160	0.145	-0.083	<0.0001

Notes: Cross-sectional comparisons are made between census block groups ranked above and below the median of selected characteristics. For each block group, total road length is the sum of highway miles within a 3-mile radius of the centroid. Vegetation is the share of land categorized as deciduous forest, evergreen forest, mixed forest, shrub/scrub, grassland/herbaceous, pasture/hay, and cultivated crops per the National Land Cover Dataset 2012. Median values are computed conditional on being within a policy city and inside the policy boundary. For road length, percentage differences are estimated using log regression. For vegetation, differences are reported in absolute levels.

Lastly, I examine whether more-affected neighborhoods also experience larger changes in resident size and find limited differences. Appendix Figure B.5 shows some, but statistically insignificant, differences in resident size by neighborhood characteristics. Within the policy limits, neighborhoods with higher shares of Hispanic residents, service-sector workers, or individuals in poverty experience larger population declines. Outside the policy limits, there is little evidence of residential sorting along these demographic lines, contradicting the conjecture that air quality impacts are driven by residential population-induced activities.

2.6.2 Changes in Spatial Distribution by Worker Types

The previous section presents some, but limited, heterogeneities in total resident size by neighborhood demographic and socioeconomic conditions. This may seem counter-intuitive given that minority, service, and less-educated workers are more likely to be lower-wage workers and thus more affected by the policy (Card & Krueger 1994; Dube et al. 2007; Cengiz et al. 2019; Dube & Lindner 2024). To resolve this tension, I leverage subgroup files in the LEHD origin–destination data and explicitly retrieve the number of workers for specific worker types. That is, I use the number of Hispanic, less-educated, or service-sector workers by residence/workplace block groups as outcome variables for this exercise. Consistent with the literature, I find larger impacts for Hispanic, service-sector, and less-educated workers, while higher-wage workers experience minimal impacts.¹⁴

Table 2.7 reports policy impacts on the spatial distribution of Hispanic, service-sector, and less-educated workers, plus a placebo group of high-wage workers. As discussed above, outcome variables represent the numbers of subgroup workers by resi-

¹⁴Less-educated workers are defined as those without a four-year college degree. Service workers are those in NAICS sector code 72 (Accommodation and Food Services). Higher-wage workers are defined as those in NAICS sectors 51 (Information), 52 (Finance and Insurance), and 54 (Professional, Scientific, and Technical Services).

dence/workplace block groups. The left column shows estimates for block groups within the policy limits, representing the policy impacts on the workers who work (employment size) or reside (resident size) in policy cities. Compared to the population average decline of -0.59% , Hispanic and less-educated workers see larger employment declines of -0.81% and -1.02% , respectively. Employment size is not significantly affected for service-sector workers. All three lower-wage groups experience larger resident-size declines (-1.07% to -1.49%), exceeding the overall -0.97% estimate. Outside the policy limits (Column 2), effects are generally smaller, though less-educated workers see a small but significant decline in employment size, indicating that areas within the policy limits become less attractive places to live and work for these workers.

These results imply that policy impacts on the spatial distribution of workers are substantially larger for Hispanic, less-educated, and service workers. In contrast, the placebo group (high-wage workers) shows increased employment size and negligible changes in resident size; these workers likely face little direct wage change from the policy and have greater income flexibility. Taken together, these results echo and extend findings on employment impacts of minimum wage policies in the literature. My results are similar to Neumark et al. (2004); Monras (2019); Dube & Lindner (2021); Jardim et al. (2022); and Wursten & Reich (2023), who identify negative employment impacts when minimum wage policies raise binding wage levels. The sharper declines among Hispanic and less-educated workers are consistent with evidence that minimum wage increases disproportionately affect groups for whom the policy is more binding (Dube & Lindner 2024; California State Legislature 2024). Moreover, the residential shifts align with Martin & Termos (2015); Monras (2019); and Dustmann et al. (2022), who show that minimum wage policies affect where jobs and workers are located.

Table 2.7: Heterogeneous Impacts on Spatial Distribution of Workers by Worker Types

	Inside Policy Limits	Outside Policy Limits
<i>Panel A. Hispanic</i>		
Employment size	-0.0081 [0.0842]	-0.0036 [0.4289]
Resident size	-0.0149* [0.0321]	-0.0072 [0.4128]
<i>Panel B. Service</i>		
Employment size	0.0002 [0.8457]	0.0030 [0.6693]
Resident size	-0.0107* [0.0120]	0.0019 [0.5852]
<i>Panel C. Less-Educated</i>		
Employment size	-0.0102 [0.0521]	-0.0054** [0.0040]
Resident size	-0.0137** [0.0040]	-0.0017 [0.5651]
<i>Panel D. "Placebo": Higher-Wage</i>		
Employment size	0.0161 [0.1122]	-0.0080 [0.2204]
Resident size	-0.0012 [0.8257]	-0.0018 [0.4890]
N	42808	80486


Notes: Estimations follow equation 2.10, separated by census block groups within (column 1) and outside (column 2) the policy boundaries. Outcomes are employment/resident size of specific worker subgroups in logarithmic form. All specifications control for block-group and county-by-year fixed effects. **Wild clustered bootstrap P-values** are reported in square brackets.

2.6.3 Residential Sorting by Amenity and Commute Accessibility

The previous subsection shows that certain workers are more likely to move away from policy cities, yet they are not sorted into neighborhoods that necessarily have the demographic characteristics described above. To explore features that moderate residential re-sorting, I integrate these findings with the conceptual model in Section 2.2 to test the model’s prediction that amenities—such as access to business establishments—and commute feasibility influence workers’ residential choices.

The left panel of Figure 2.5 shows that workers are relocating to neighborhoods outside the policy limits with better non-environmental amenities. Here, a better non-environmental amenity is defined as greater ease of access to nearby business establishments, following Cook (2025). This figure presents policy impacts on resident size using the same empirical strategy as in Section 2.6.1. Within the policy limits, low-amenity block groups experience a large loss in resident size, while the high-amenity group sees a smaller change. This suggests that low-amenity neighborhoods are less able to retain residents when faced with rising living costs. Outside the policy limits, high-amenity neighborhoods gain residents while low-amenity neighborhoods lose them, although the difference is not always statistically significant.

Figure 2.5: Effects on Workers' Residence by Neighborhoods' Location and Characteristics



Notes: The left panel reports policy impacts on residential population size by block group location $\times$ neighborhood amenity index (top/bottom 50%). The Amenity Index is sourced from the Neighborhood Amenity Quality Index (Cook 2025) and is based on visits to establishments made by low-income households. A higher value indicates better access to neighborhood amenities. The right panel reports the same estimates by commuting accessibility, measured using the pre-policy median commuting time from the 2010–2014 5-year ACS. Within each panel, estimates are presented (from left to right) for block groups that are: (1) inside policy limits and above-median; (2) inside policy limits and below-median; (3) outside policy limits and above-median; and (4) outside policy limits and below-median. Outcome variables are in logarithmic form. Standard errors are clustered at the policy city level.

The right panel presents analogous estimates by pre-policy commuting-time ranking. Within the policy limits, neighborhoods with shorter pre-policy commutes experience

marginally greater residential losses. Outside the policy limits, workers appear to relocate to neighborhoods with shorter commutes and away from those with longer ones. This finding reflects cost-minimization behavior: workers seek to reduce combined living and commuting costs by relocating to neighboring areas with better access to employment. Differences on this dimension are not statistically significant.

Appendix Figure ?? examines whether this sorting leads to differential congestion and air quality impacts based on amenity and commuting time. The results show modest differences, with larger increases in traffic and pollutant concentrations in low-amenity neighborhoods outside the policy limits. These results, together with the sorting patterns in Figure 2.5, indicate that traffic impacts are unlikely to be driven primarily by increases in residential population. More specifically, the divergences between population growth and air quality suggest that household emissions—such as those from residential cooking and heating—are unlikely to be the primary drivers of the observed air quality changes. On the other hand, neighborhoods with worse access to business establishments also tend to have worse green amenities: they have about 11.1% more roads and a 0.096 lower vegetation share. The correlation between environmental amenities (e.g., green space) and non-environmental amenities (e.g., access to establishments) can help explain why these neighborhoods are more affected.

2.7 Policy Simulation

This section evaluates the aggregate impact of city-level minimum wage ordinances. I first integrate a traffic-emissions model with a particle dispersion model to link traffic impacts directly to pollutant concentrations using a two-stage framework. Next, to better understand the magnitude of policy impacts and to provide policy insights, I present comparable simulation results for a 10% change in gasoline prices. Finally, I analyze

the marginal net benefit of city-level minimum wage ordinances for lower-wage workers, accounting for congestion and air quality consequences.

2.7.1 Traffic-Related Emissions and Pollutant Concentrations

Previous sections reported the policy’s impacts on traffic and air quality separately. To establish a direct link between these two outcomes and isolate the portion of air quality change attributable solely to traffic, I simulate and compare traffic-related emissions and pollutant concentrations under the observed scenario with a counterfactual scenario in which the city-level minimum wage ordinances are absent. In other words, the average counterfactual minimum wage is represented by the bottom dashed line in Figure 2.2.


I use a two-stage modeling approach, combining the EMISSION FACTORS (EMFAC) model and the Intervention Model for Air Pollution (InMAP) to simulate vehicle emissions and the subsequent dispersion of pollutants and their effect on ambient concentrations. For the benchmark scenario, I input observed traffic data into EMFAC to obtain emissions of multiple pollutants, including NO_x , $PM_{2.5}$, SO_x , volatile organic compounds, PM_{10} , CO , and CO_2 . These emissions are then fed into InMAP to simulate the resulting pollutant concentrations and $PM_{2.5}$ -related mortality under the prevailing policy environment. To improve precision and avoid spatial mismatch, I re-estimate the effect of minimum wage increases on traffic at the road-segment level, allowing impacts to vary by hour of the day and day of the week.

To construct the counterfactual (no-policy) scenario, I calculate the difference between state and city-level minimum wages, ΔMW , which reflects the marginal contribution of city-level ordinances.¹⁵ I subtract the estimated policy-induced changes (the product of estimated coefficients and ΔMW) from observed traffic data to construct counterfactual traffic measures. These adjusted values are then input into EMFAC to estimate coun-

¹⁵See Figure 2.2 for the average counterfactual minimum wage.

terfactual emissions, which are fed into InMAP to simulate pollutant concentrations and $PM_{2.5}$ -related mortality in the absence of city-level wage policies. I also report the change in total vehicle miles traveled and calculate the corresponding change in traffic fatalities using fatality rates from the Department of Transportation. Figure 2.6 summarizes the simulation process.

Figure 2.6: Simulation Process



Notes: This diagram outlines the simulation procedure used to generate Tables 2.8 and 2.9. I simulate real-world emissions, pollutant concentrations, and $PM_{2.5}$ -related mortality using EMFAC and InMAP, and compare these results with hypothetical scenarios based on previous estimates or literature.

Table 2.8 reports percentage changes in traffic-related emissions and pollutant concentrations, along with additional annual deaths between the observed and counterfactual scenarios. The top panel presents changes in emissions and vehicle miles traveled at the road-segment level; the bottom panel summarizes changes in pollutant concentrations and annual deaths at the census block group level.

Table 2.8: Policy-Induced Traffic-Related Emissions, Pollutant Concentrations, and Mortality

<i>Panel A: Road Segment Level</i>					
<i>Traffic-Related Emissions (%)</i>					
	Total	High VMT	High Emission	High ΔMW	
NO_x	2.328	1.384	2.097	5.292	
$PM_{2.5}$	2.785	1.660	2.449	5.517	
PM_{10}	2.780	1.657	2.445	5.516	
SO_x	1.592	1.075	1.545	4.794	
CO	1.228	0.829	1.136	4.794	
CO_2	1.600	1.077	1.552	4.800	
<i>Annual Extra VMT and Traffic Fatality</i>					
Vehicle Miles Traveled	1.16×10^9	4.08×10^8	4.95×10^8	6.85×10^8	
Traffic Fatalities	13.161	4.605	5.595	7.745	
<i>Panel B: Block Group Level</i>					
<i>Traffic-Related Pollutant Concentrations (%)</i>					
	Total	Most Polluted	Second Most Polluted	Most Affected	
NO_x	0.446	1.477	0.733	3.530	
SO_x	0.457	1.458	0.741	3.472	
$PM_{2.5}$	0.451	1.481	0.737	3.538	
<i>$PM_{2.5}$-Related Mortality (Deaths/Year)</i>					
Total	1.958	1.246	0.399	1.281	
Hispanics	0.591	0.362	0.138	0.315	
White	0.738	0.461	0.131	0.477	
Asian	0.309	0.203	0.059	0.239	

Notes: This table reports percentage changes in traffic-related emissions and pollutant concentrations between the observed scenario and a counterfactual case where traffic effects of minimum wage ordinances are removed. To improve precision, traffic effects are re-estimated at the road-segment level without weighting, using post-policy observations only. Emissions are simulated using EMFAC, based on total vehicle miles and average speed. Pollutant concentrations and annual $PM_{2.5}$ -related mortality are then derived using InMAP. The first column reports results for the full sample. Three sub-samples of block groups (road segments) are also presented. In the top panel, “High VMT” and “High Emission” refer to segments in the top 20% of pre-policy vehicle miles and pollutant emissions, respectively. In the bottom panel, “Most Polluted” and “Second Most Polluted” refer to block groups in the top two quintiles of pre-policy $PM_{2.5}$ concentrations. The final column in both panels reports estimates for block groups with the highest ΔMW .

Column 1 reports results for the full sample, while the next three columns show distributional effects for subgroups. The top panel shows that, considering only traffic impacts, city-level wage ordinances increase pollutant emissions by 1.228% to 2.785

The bottom panel reports changes in pollutant concentrations and associated $PM_{2.5}$ -related deaths within the sample geography. Column 1 shows that minimum wage increases raise NO_x concentrations by 0.446%, SO_x by 0.457%, and $PM_{2.5}$ by 0.451%. These increases result in an average of 1.958 additional annual deaths (valued at \$29.83 million based on the EPA’s value of a statistical life). Consistent with earlier findings, air quality effects are concentrated in already polluted areas. Block groups in the top quintile of pre-policy traffic-related $PM_{2.5}$ concentrations experience a 1.481% increase, leading to 1.246 additional deaths annually. The next quintile experiences a 0.737% increase and 0.399 additional deaths. In areas with the top 20% ΔMW , $PM_{2.5}$ rises by about 3.5%, with 1.281 additional deaths per year. These health cost estimates are conservative, as they consider only $PM_{2.5}$ -related mortality, excluding other health costs and impacts from additional pollutants.

2.7.2 Policy Contextualization: A 10% Change in Gasoline Prices

While Section 2.7.1 illustrates the air quality implications of wage ordinances, it is important to benchmark these effects against other policy instruments. I therefore conduct a similar simulation for a 10% gasoline tax (approximately \$0.372 per gallon). Drawing from Mamkhezri & Khezri (2024), I model VMT as a function of gasoline prices, using their elasticity estimate of -0.257 to capture the traffic response to a 10% price increase for both policy cities and neighboring block groups. I then apply the same emissions and air quality models to estimate the resulting changes.

Table 2.9 summarizes percentage changes in emissions, air quality, and $PM_{2.5}$ -related

deaths under the gasoline price change, relative to the baseline scenario. To simplify the comparison and highlight magnitudes, values in the table are also expressed as positive numbers. Results indicate that the tax changes emissions by 0.974% to 1.610%, which is about 60% of the emissions generated by the minimum wage policy. For vehicle miles traveled, the total change is 9.26×10^8 , associated with 10.467 additional traffic fatalities. Corresponding pollutant concentrations change by up to 0.457%, accounting for at least 91.9% of the increase caused by local wage floors. Mortality increases by 1.722 deaths per year within the sample geography, accounting for 87.8% of the excess deaths attributed to the minimum wage ordinances.

Table 2.9: Changes in Traffic-Related Emissions, Pollutant Concentrations, and $PM_{2.5}$ Mortality with a 10% Change in Gasoline Prices

<i>Panel A: Road Segment Level</i>	Total	High VMT	High Emission
<i>Traffic-Related Emissions (%)</i>			
NO_x	1.419	1.086	1.404
$PM_{2.5}$	1.610	1.263	1.591
PM_{10}	1.608	1.261	1.589
SO_x	1.149	0.943	1.176
CO	0.974	0.809	0.976
CO_2	1.150	0.941	1.176
<i>Annual Extra VMT and Traffic Fatalities</i>			
Vehicle Miles Traveled	9.26×10^8	3.84×10^8	4.12×10^8
Traffic Fatalities	10.467	4.342	4.659
<i>Panel B: Block Group Level</i>	Total	Most Polluted	Second Most Polluted
<i>Traffic-Related Pollutant Concentrations (%)</i>			
NO_x	0.410	1.308	0.572
SO_x	0.457	1.342	0.597
$PM_{2.5}$	0.409	1.312	0.574
<i>$PM_{2.5}$-Related Mortality (Deaths/Year)</i>			
Total	1.722	1.144	0.316
Hispanics	0.554	0.348	0.119
White	0.647	0.430	0.101
Asian	0.263	0.181	0.046

Notes: This table reports changes (in absolute values) in emissions and air quality levels from the benchmark to a counterfactual scenario in which gasoline is taxed at 10%, using price elasticity estimates from Mamkhezri & Khezri (2024). Emissions, pollutant concentrations, and mortality are simulated sequentially using the EMISSION FACTORS (EMFAC) model and the Intervention Model for Air Pollution (InMAP). The first column reports estimates for the entire sample in both panels, while the two sub-samples of block groups follow definitions provided in Table 2.8 and Section 2.7.1. For consistency, the results cover the same period as Table 2.8.

2.7.3 Cost-Benefit Analysis

The primary motivation for city-level minimum wage ordinances is to improve the living standards of lower-wage workers. However, the results in Section 2.5 suggest

that higher rental prices, increased congestion, and worsened air quality offset much of the wage benefit. To quantify the net marginal impact of city-level minimum wage ordinances, I combine estimates from this paper with price estimates from the literature to compute the marginal monetary gain for a representative worker following a \$1 increase in the minimum wage.

Table 2.10 summarizes the marginal costs and benefits for a representative low-wage worker who works within city limits but commutes from neighboring areas, facing a \$1 increase in the city-level minimum wage. Earnings and price estimates are drawn from Leung (2021) and Jardim et al. (2022), while rent, food prices, commute time, congestion, and air quality estimates come from Section 2.5. The policy raises quarterly earnings by \$76.60 for low-wage workers but imposes costs through higher rent and food prices, longer travel times due to congestion, and elevated health risks. Among these, congestion accounts for the largest offset (around 30%), exceeding the combined effects of higher food and rental costs.

Table 2.10: Policy Evaluation: Annual Marginal Costs and Gains for a Low-Wage Worker

Channel	Estimates	Source	Other Assumptions	Δ \$
Earnings	Quarterly \$76.6 per \$MW	Jardim et al. (2022)	-	+306.40
Rental	0.085% per \$MW	Table 2.4	Monthly rents \$1270	-12.95
Food	0.8% per 10% MW	Leung (2021)	Annual food expenditure \$8629	-66.31
Traffic	+0.0736 minutes; +0.0691 miles; and 3.84% more congested per \$MW	Tables 2.5, 2.3	2-way per workday, wage at \$14 VSL at 8.7 million; 1.13 deaths per 100 million miles driven	-106.54
Health	0.74% in $PM_{2.5}$ 0.61% in NO_2 per \$MW	Table 2.3	VSL at 8.7 million	-14.15
Net Gain				+106.45

2.8 Conclusion

This paper demonstrates that place-based minimum wage ordinances, while intended to improve the welfare of low-wage workers, can generate significant, negative environmental externalities. Through a comprehensive analysis of traffic data, satellite-based air quality measures, and worker commuting patterns, I show that these policies worsen traffic congestion and air quality in areas surrounding policy cities. The findings reveal a disproportionate burden borne by disadvantaged neighborhoods and emphasize the importance of accounting for spatial spillovers and unintended consequences in the design of place-based policies.

Furthermore, the study contributes to the understanding of intra-urban residential sorting by showing how wage increases at the bottom of the distribution influence location decisions. The simulation results underscore the trade-offs between labor market benefits and environmental costs, suggesting that complementary measures—such as investments in public transportation or green infrastructure—may be necessary to ensure that labor market interventions do not come at the expense of environmental justice and public health. These results point to a broader lesson: the success of place-based policies hinges not only on their direct economic effects but also on how they shape behavior, economic activity, and environmental conditions. Policymakers should integrate economic and environmental planning to mitigate unintended outcomes and enhance overall effectiveness of such policies.

Chapter 3

How Did Safety-Net Reform Affect the Education of Adolescents from Low-Income Families?

3.1 Introduction

The U.S. safety net was extensively revamped in the 1990s. After a period of experimentation by the states, federal welfare reform replaced the cash transfer program, Aid to Families with Dependent Children (AFDC), with the Temporary Assistance to Needy Family (TANF) program, in 1996. Congress expanded the Earned Income Tax Credit (EITC) over the period 1994-1996, raising wage subsidies available to low-wage workers. Changes to the Medicaid program, which began in the 1980s, continued during the 1990s to extend health coverage to children in low-income families.

An important question is, how did these reforms affect the educational attainment of individuals who were adolescents at the time the reforms were enacted? This question is important for policy purposes, since the goal of policy makers was to change the safety

net in ways that would have long-term effects on U.S. society. It may also have important implications for trends in inequality.

Several authors have analyzed how specific reforms have affected the schooling of individuals who were adolescents at the time when the reforms were passed. Dave et al. (2012) and Kaestner et al. (2003) analyzed how welfare reform affected dropout behavior. Maxfield (2013), Bastian & Micheltore (2018), and Manoli & Turner (2018) estimated the effect of EITC expansions on educational attainment. The effect of Medicaid expansions on college attendance has been studied by Cohodes et al. (2016), Miller & Wherry (2019), and Brown et al. (2020).

Whereas previous authors have sought to estimate the effect of specific programs, our goal in this paper is to estimate their joint effect, that is, their effect as a bundle. One reason is that we might expect the effects of those programs to interact. For example, a healthier adolescent may be better able to take advantage of additional parental expenditures on her schooling.

Another reason is that there are at least two channels by which the reforms could have affected individuals who were adolescents at the time they were put into place. The first we refer to as the resource channel. This includes the effects of the reforms on program services, parental time spent with children, and parental income, which in turn could affect the adolescent's ability to afford post-secondary schooling. The resource channel has been the implicit focus of previous work, as we discuss below. We would expect males and females to be similarly exposed to the resource channel, since they grow up in the same families.

However, safety-net reform may have also affected those adolescents via their expectations about the future, including their expectations regarding future safety-net benefits and eligibility. Heckman et al. (2002) present a model that predicts that early-career wage subsidies may reduce human capital investments that are rivalrous with work. Thus ex-

expectations of future wage subsidies could reduce schooling. Studies that focus on the resource channel may miss the expectations channel. In contrast to the resource channel, we would expect the expectations channel to affect males and females differently, since in the pre-reform regime, women were much more likely than men to receive safety-net benefits.

We take a difference-in-differences (DD) approach to estimation, comparing two groups across two generations of adolescents. The pre-reform generation consists of the 1961-1964 birth cohorts of the 1979 National Longitudinal Survey of Youth (NLSY79). This generation was aged 11-15 in 1975, and in their 30s by the time of safety-net reform. The post-reform generation comes from the 1997 National Longitudinal Survey of Youth (NLSY97), whose members were born between 1980 and 1984. This generation was aged 10-15 in 1995, in the midst of the 1994-96 EITC expansions and just before federal welfare reform. Members of this generation would have been affected to varying degrees by the Medicaid expansions, as well.

Within each generation, we define a treatment and a comparison group. The treatment group consists of individuals whom we predict to be at high risk of economic disadvantage during adolescence. We show that this group has characteristics similar to those of groups at high risk of welfare receipt in the mid 1990s. As a result, they should have been affected by safety-net reform. The comparison group consists of individuals at lower risk of adolescent disadvantage, who would have been less affected by safety-net reform.

The logic of differences-in-differences here is straightforward. The contrast between the high- and lower-risk groups in the pre-reform generation establishes the baseline difference in their educational attainment prior to safety-net reform. Differencing across generations nets out changes over time in factors that influence schooling which affect both groups similarly. If safety-net reform were the only factor that affected the two

groups differently, then the DD approach would identify the bundled effect of the various safety-net reforms. This would include the effects of any interactions as well as effects that arise via resource and expectations channels.

Of course, other factors may have affected schooling trends as well. These include differences in the U.S. economy as faced by the two generations as they finished high school. We expand on the basic difference-in-difference approach by adding various macro- and socio-economic controls to account for these factors.

Our focus here is on educational attainment at age 25. We consider years of schooling, high school graduation, sub-baccalaureate degrees, and college graduation. We find that safety-net reform may have reduced educational attainment for women, while increasing it for men. Women obtained fewer years of schooling, were less likely to graduate high school, and less likely to obtain a post-secondary degree. However, although all these effects were in the same direction, they were marginally significant. In contrast, men were significantly more likely to obtain a post-secondary degree.

To the best of our knowledge, these results provide the first estimates of the bundled effects of safety-net reform on schooling. Our findings for women differ from those of some earlier studies, which focused on the resource channel for specific safety-net programs. To help explain these differences, we analyze annual enrollment and labor supply between the ages of 18 and 22, when most people who attend college would be enrolled. We find that young women are less likely to be enrolled and more likely to work during this period as a results of safety-net reform. This is particularly true for young women with children. As we explain below, this may be consistent with the notion that the expectations channel led to the reduction in human capital among women. In a follow-up paper Bastian et al. forthcoming, we build on these results, and examine how these changes in human capital affected employment, safety-net participation, marriage, and childbearing up through age 32.

In the next section of the paper, we provide more background to motivate our question and briefly discuss the various safety-net reforms. After discussing the data in the following section, we elaborate our empirical approach. A key to our DD analysis involves constructing our treatment and comparison groups, which in turn involves predicting the risk of adolescent disadvantage among the members of our sample. We discuss how we predict risk and how we construct the risk groups. The approach involves a number of decisions whose effect on our results is impossible to gauge *ex ante*. We discuss the specification-curve approach to inference of Simonsohn et al. (2020), which allows us to draw inferences across the multiple specifications that result from those decisions. Section 4 also presents our results, which show that safety-net reform increased men’s schooling, although it may have reduced it among women. It also provides evidence on employment and enrollment behavior. Section 5 discusses the role of resource and expectations channels and offers suggestions for why our results run contrary to recent studies that have analyzed individual components of the safety net.

3.2 Background

Here we discuss the various safety-net reforms with an eye to two key questions: (1) how should the reforms have affected the resources available to individuals who were adolescents at the time they took place? and (2) how might the reforms have affected their expectations? While several aspects of safety-net reform increased household earnings, the overall effect of safety-net reform on household resources and adolescents’ expectations varied quite a bit, and is not obvious *ex ante*.

In our discussion of the recent literature, we focus solely on studies that have analyzed the link between safety-net reforms and educational attainment of children who were in school when the reforms took place. In Bastian et al. (forthcoming), we discuss additional

policy-change details and studies that analyze later-life outcomes.

3.2.1 Welfare reform

Before the 1990s, welfare (i.e., AFDC) provided cash aid to poor families with children. AFDC's work requirements were lax, and any earned income declared by recipients was taxed away, in the form of reduced benefits, at a high rate. In the 1990s, states began to experiment with their AFDC programs and introduce restrictions. Roughly half the states strengthened their work requirements, stiffened sanctions for non-compliance, reduced the implicit tax rate on earnings, and/or imposed time limits. In 1996, federal legislation replaced AFDC with Temporary Assistance for Needy Families (TANF), which incorporated many of these changes nationwide.

These reforms should have affected resources available to adolescents in families receiving aid. All would be expected to increase parental employment, and thus to reduce the amount of time that the adolescent's parent spent in the home. How the reforms should have affected total family income depends on whether earnings rose by more than welfare benefits fell.

The expectations channel of welfare reform would affect females more than males, since about 90 percent of adult recipients of welfare were single mothers Services & Human 1995. One might expect that the reduced generosity of welfare, stemming from the combination of time limits and work requirements, would lead adolescents at the time to expect to work for a greater share of their adult lives. This should raise the value of human capital, which could take the form either of schooling or of labor market experience.

A sizeable empirical literature, summarized by Blank (2002), Grogger & Karoly (2005), Moffitt (2003), and Ziliak (2008), found that welfare reform indeed increased

employment among single mothers. Effects on household income were less clear, with some studies finding largely null effects Bavier 2001, US Department of Health and Human Services 2001.

Roughly two dozen experimental studies analyzed the effects of various state-level welfare reforms on the children of recipients (see Grogger & Karoly (2005) for a summary). The results were quite heterogeneous and many estimates were insignificant. Those results which were significant suggested that younger children of recipients may have benefited from the reforms. However, for individuals who were adolescents at the time of reform, there is some evidence of negative effects on school performance Gennetian et al. 2002 and higher rates of risky behavior, including truancy and delinquency Dave et al. 2020. Both of these studies focus on adolescents who were contemporaries of the members of our post-reform generation.

Two studies have focused on high school dropout among individuals who were adolescents at the time of reform. Kaestner et al. (2003) used a design similar to ours, following groups at varying levels of risk in the NLSY79 and NLSY97. Their high-risk group consisted of adolescents in families with single parents and/or parents with low levels of schooling, whereas the low-risk group consisted of everyone else. Their results showed that the high-risk group in the latter generation was less likely to drop out of high school at ages 17 and 19. However, the results were stronger at age 17 than at age 19, consistent with a provision of federal welfare reform that required minor children to be enrolled in school in order for the family to receive aid. Thus their results may indicate primarily that welfare reform succeeded in reducing early dropout.¹

Dave et al. (2012) studied high school dropout among 15- to 19-year-olds using data from the CPS. To estimate the effect of welfare reform, they made use of differences in

¹When we employ their sample and their definitions of risk groups, we too find reductions in high school dropout at ages 17 and 19. If we instead use outcomes measured at age 25, we obtain results similar to those reported below.

the timing of reforms across states. Their results showed a reduction in dropout due to reform, but considering the age of their sample, it is again hard to say whether this represents a reduction in early dropout or a true increase in educational attainment. An important difference between both of these studies and ours is that we study educational attainment at age 25, which for the vast majority of the population represents completed schooling.

3.2.2 The EITC

The EITC provides a refundable tax credit that varies according to the family's earnings and number of children. In 1993, Congress legislated expansions to the EITC that were phased in between 1994 and 1996. They increased the generosity of the program and greatly increased the relative generosity towards families with more children. Between 1993 and 1996, the maximum benefits rose from \$1434 to \$2150 for one-child families, and rose from \$1511 to \$3556 for families with two or more children Hotz & Scholz 2003.

Basic labor supply theory makes several predictions regarding the effect of the EITC on the resources available to children. It predicts that the EITC should increase the employment of single parents along the extensive margin, presumably reducing the time they spend with their children Bastian & Lochner 2020. Whether the EITC increases family income depends on whether any earnings gains exceed any losses in non-wage benefits.

An extensive literature shows that the EITC increases the employment of single mothers, who account for roughly 60 percent of EITC recipients Hotz & Scholz 2003, Nichols & Rothstein 2016, Hoynes & Patel 2018, Bastian 2020.² Several studies show negative, but generally small, reductions in labor supply among secondary earners in two-parent households Eissa & Hoynes 2004. Hoynes & Patel (2018) and Bastian & Jones (2020)

²Kleven (2019) presents results to the contrary.

show that the EITC raises after-tax-and-transfer income among single mothers.

Three recent studies linking the EITC to higher educational attainment among individuals whose families received the EITC while they were children. Maxfield (2013) and Bastian & Micheltore (2018) exploit variation in the generosity of the EITC over time and by family size to estimate the relationship between EITC payments received by a family and the eventual educational attainment of its children. Manoli & Turner (2018) also finds that the EITC increases college attendance for youth in their last year of high school. By exploiting variation in EITC benefits across families and over time, while the children were young, these studies clearly focus on the resource channel.

As mentioned above, the EITC may not only affect resources available in the parental home, it may also affect the adolescent's expectations regarding labor market rewards. Heckman et al. (2002) point out that wage subsidies during the schooling years promote employment and raise the opportunity cost of schooling. Education may fall as a result. They provide some simulations, based on CPS data, that are broadly consistent with this prediction. However, we are not aware of any direct tests for substitution between schooling and work as a result of the EITC.

3.2.3 Medicaid

Prior to 1984, Medicaid coverage for poor families with children was limited to families receiving AFDC. Between 1984 and 1997, it was expanded several times. Some of the expansions applied to specific birth cohorts of children and/or applied to children below certain ages. The expansions also had different income-eligibility thresholds and varied in whether states were required to expand coverage or given the option to do so.

An abundant literature, surveyed by Gruber (2003) and Buchmueller et al. (2016), shows that the various expansions greatly increased resources, in the form of health

coverage, to children in poor families. Likewise, numerous analysts have established a link between Medicaid eligibility during childhood and favorable childhood outcomes Buchmueller et al. 2016.

Most relevant for our work are studies that relate childhood Medicaid expansions to educational attainment. Miller & Wherry (2019) find that perinatal health coverage has the long-run effect of increasing high school graduation rates. Brown et al. (2020) show that each additional year of Medicaid eligibility during childhood increases the probability of having ever enrolled in college. Cohodes et al. (2016) find that increases in Medicaid eligibility during childhood decreases high school dropout rates, increases college enrollment, and increases college completion.³

Beyond these effects via the resource channel, Medicaid may also have an effect via the expectations channel. Low-wage jobs often fail to provide health insurance, either for the worker or her children. Higher-wage jobs usually provide such benefits, but require higher levels of education. By weakening the link between schooling and health coverage for one's children, the Medicaid expansions may reduce the return to college attendance, in turn reducing investment in this form of human capital. However, we are aware of no empirical studies on this matter.

3.2.4 Other safety net reforms

Two other programs which contribute to the safety net have recently been reformed. The Supplemental Nutritional Assistance Program (SNAP; formerly Food Stamps) was expanded in response to the Great Recession Ganong & Liebman 2018. We do not emphasize these expansions here because most of them took place after our sample members completed their schooling. The Child Tax Credit (CTC) was enacted in 1997. Prior to

³These findings are broadly consistent with those of Boudreaux et al. (2016) and Goodman-Bacon (2016), who studied the long-term effects of the initial roll-out of the Medicaid program.

2001, it was non-refundable, limiting its value to low-income families. In 2009, substantial benefits became available for low-income families Chetty et al. 2011, Nichols & Rothstein 2016, Goldin & Micheltore 2020. By then, our sample members had largely finished their education.

3.3 Data

3.3.1 Survey basics

The NLSY comprises two longitudinal surveys that have followed two generations of Americans over a lengthy period of time. The NLSY79 began in 1979 with 12,686 individuals who were born between 1957 and 1964. They were interviewed annually between 1979 and 1994 and biennially between 1994 and 2016. As we discuss below, various subsamples have been discontinued over time, but the retention rate for the remainder of the sample ranged from 96 to 84 percent between 1980 and 1998, which are the interviews that we make use of here.

The NLSY97 began in 1997 with 8,984 individuals who were born between 1980 and 1984. They were interviewed annually between 1997 and 2011 and biennially between 2011 and 2017. In 2017, the retention rate was 75 percent. Both surveys are based on stratified probability samples of the US population in their respective birth cohorts.

3.3.2 Sample exclusions

To make the samples as comparable as possible, we restrict attention to particular subsets of survey respondents. The NLSY79 initially included a military sample as well as so-called oversamples of low-income youths. The military sample was largely discontinued in 1985 and the white low-income oversample was discontinued in 1991. We exclude both

of these groups from all of our analyses.

We make other exclusions based on age. Certain variables from “wave 0,” that is, the year prior to the first interview, are available only for those respondents who lived with a parent figure at the time of the first interview. Therefore, we limit our main analyses to NLSY79 respondents who were born between 1961 and 1964, and were therefore 14-17 years old in 1978. In the NLSY97, respondents were aged 12-16 in 1996, so we make no age-related exclusions.

Finally, we drop Hispanic respondents from both surveys. The reason is that the Hispanic population of the United States was very different in 1979 and 1997, rendering the Hispanic subsamples of the two surveys largely non-comparable. The 1979 survey began when the United States was near a long-time low in immigration. In contrast, the 1997 survey began after immigration had been rising for nearly two decades, with substantial numbers coming from Spanish-speaking countries Migration Policy Institute 2019. Hispanics in the NLSY79 were likely members of families who had been in the United States for a considerable time, whereas many Hispanics in the NLSY97 were first-generation immigrants Duncan et al. 2020.

3.3.3 Key variables

Key explanatory variables include the mother’s education, father’s education, family structure at age 12 (NLSY97) or 14 (NLSY79), and family size, all of which were collected in the first interview. Also collected at that time was information on household income and whether the household had received cash welfare during the previous year. To account for the potential effects of macro- and socio-economic factors on schooling, we merged annual state-level information on unemployment, GDP/capita, returns to schooling, income inequality, and incarceration rates onto the individual-level records.⁴

⁴The Online Data Appendix provides sources and methods for these variables.

We then averaged the values of those variables over the years in which the adolescent was 18-22 years old.

Our analysis makes use of two sets of dependent variables: a measure of economic status during adolescence, which we use to construct our treatment and comparison groups, and several measures of educational attainment. Economic status during adolescence is measured as the ratio of income to needs. Specifically, it is constructed as the average value of household income relative to the household's poverty threshold when the sample member was between the ages of 14 and 17.⁵ The idea is that adolescents from families with lower income/needs would be at greater risk of being poor as adults themselves, and therefore be more likely to be affected by changes to the safety net. Adolescents at lower risk would be less likely to be affected by safety net reform.

We analyze several measures of educational human capital. The first is highest grade completed. We add to that three measures of degree attainment. The first is an indicator equal to one if the youth received a high school credential, including a GED, and equal to zero otherwise.⁶ The second is an indicator of whether the youth received a post-secondary credential that was lower than a bachelor's degree. The final one is an indicator of whether the youth received a bachelor's degree. We refer to these measures as recording educational outcomes at age 25. For about 90 percent of the sample, that is literally true. In order to minimize missing values, for youths who were missing data at age 25, we filled in with data from nearby ages.⁷ Sample means for all the variables that appear in our analysis are shown in online Appendix Table 1.

⁵The NLSY surveys contain many households with multiple sibling respondents. We average household income and poverty thresholds across all siblings over the relevant age ranges

⁶The GED does not factor into the calculation of highest grade completed.

⁷For youths who were missing age 25, we looked at age 26. For youths missing ages 25 and 26, we looked at age 24, and likewise for ages 27 and 23.

3.4 Estimation and Results

3.4.1 Constructing treatment and comparison groups

Risk regressions and risk groups

To construct our treatment and comparison groups, we first estimate a regression to predict income/needs during adolescence. We use the predicted values from that regression as an index of risk, then construct high- and lower-risk groups based on terciles of that index. The risk regression takes the form

$$R_{gi} = Z_{gi}\theta + u_{gi} \quad g = 1979, 1997; i = 1, \dots, n_g \quad (3.1)$$

where R_{gi} equals income/needs between the ages of 14 and 17; Z_{gi} is a vector of predictors including i 's mother's education, father's education (both linear and quadratic terms), family structure at age 12/14, sex, race, and all interactions among those variables. The term θ is a vector of coefficients to be estimated, u_{gi} is an unobservable disturbance term, and n_g is the number of observations in generation g .

Equation (1) is estimated using data from the NLSY79. We then use the resulting regression coefficients to predict adolescent income/needs for both the NLSY79 and the NLSY97, on the basis of the each generation's predictor variables.⁸ We refer to these predicted values as a risk index. We then define treatment and comparison groups on the basis of the distribution of the risk index. We take the first tercile, that is, the lowest predicted income/needs, as the treatment group and the 2nd tercile as the com-

⁸In preliminary work, we also included the macro- and socio-economic variables described above in the risk regressions. However, this presented us with a serious out-of-sample prediction problem, since many of those variables exhibited strong trends between the two generations. As a result, some of the risk groups in the NLSY97 that were based on the NLSY79 terciles were nearly empty. We therefore estimated the risk regressions as a function of person- and household-level variables only. We use the macro- and socio-economic variables as controls in some of the outcome regressions, as discussed below.

parison group. We do not use the 3rd tercile for our main analysis. We show below that adolescents in the top tercile of the income/needs distribution come from much more advantaged backgrounds than those in the bottom two terciles. As a result, they may provide a poor comparison group.

To this point, we have sidestepped the fact that our approach requires us to make several decisions. One is whether to construct R_{gi} as the average of the ratio of income to needs, or as the ratio of average income to average needs? Another is which estimator to use? Another is whether we should include family structure as a regressor, since family structure in the later generation may have been affected by state-level welfare reforms in the early 1990s. Another is how to estimate the quantiles used to divide the sample into risk groups? Should we estimate them from the generation that was used to estimate the risk regression, then apply them to both generations? Or should we tabulate generation-specific quantiles instead?

Since any of these decisions could in principle affect our results, and none is obviously right or wrong, we prefer not to take a stand. Instead, we construct sixteen sets of risk groups. We estimate two versions of equation (1), one whose dependent variable is the average of the annual ratios of income/needs, and one whose dependent variable is average income divided by average needs. In each case, we use two different estimators. These include ordinary least squares and weighted least squares, weighting by the first-round sampling weights. Likewise, we estimate two sets of risk regressions, one of which includes family structure and its associated interactions, and one of which excludes them.

Finally, we use both absolute and relative terciles to construct the risk groups. In the case of absolute quantiles, we use the terciles from the NLSY79 to define risk groups for both generations. In the case of relative quantiles, we estimate separate terciles for each survey. The absolute quantiles provide an absolute definition of risk. However, changes in the distributions of parental education and family structure may cause this

approach to misclassify sample members in the latter generation. Parental education rose substantially between generations, as shown below. This means that parents who were high school graduates were probably more favorably selected in the earlier generation than the latter, which would tend to push latter-generation sample members lower in the risk distribution. How this would affect our estimates is not entirely clear. The reason is that some sample members who would have been classified as high-risk may be classified as lower-risk, but at the same time, some who would have been classified as lower-risk may be classified as lowest-risk and thus dropped from the outcome regressions altogether. Changes in the distribution of family structure give rise to similar issues. For this reason, we use the relative quantiles to provide an alternative, relative, definition of the risk groups. We combine the estimates and base inference on them using the specification curve approach of Simonsohn et al. 2020, as discussed below.

Characteristics of the high- and lower-risk groups

For the moment, we consider the results of our procedure for predicting poverty risk and classifying sample members into treatment and comparison groups. Since the coefficients from the risk regressions are rather numerous and not particularly enlightening, we present them in Appendix Table 2. In Table 3.1, we present variable means by risk group, sex, and generation. Panel A presents results for the NLSY79 and panel B presents results for the NLSY97. Statistics for women appear on the left and those for men appear on the right. In each panel, the first column reports means for the high-risk treatment group, which is the first tercile of the risk distribution, and the second column presents means for the lower-risk comparison group, which is the second tercile. For purposes of comparison, we also present means for the top tercile, in the third column, although we do not use this group in estimation below. These values are averaged over all sixteen sets of risk groups.

Many of the differences between risk groups are what one might expect. In panel A, we see that parental education is lower for youths in the high-risk group than in the lower-risk group. The high-risk youths also come from larger families, and their families had lower income at wave 0.⁹

The next three rows of the table show that our method for constructing risk groups results in a high-risk treatment group that was likely to have been directly affected by safety-net reform. First, over 30 percent of this group's families received cash welfare in the year before the NLSY79 survey began. This is similar to the level of welfare receipt at the time among single-parent families, who have often been the focus of studies of welfare incentives (e.g. Moffitt (1983), Meyer & Rosenbaum (2001), Grogger (2003)). Second, roughly 65 percent of this group was Black. Among white and black children receiving AFDC prior to safety-net reform, roughly 55 percent were black (US DHHS 1995). Finally, about 60 percent of this group lived in a single-parent household. About 90 percent of adult recipients of AFDC were single mothers (US DHHS 1995). In contrast, the lower-risk comparison group was much less likely to have received welfare in wave 0, to be Black, or to have been raised by a single parent. The high-risk treatment group was considerably more likely to have been exposed to the safety net, and thus to safety-net reform, than the lower-risk comparison group. Put differently, safety-net reform was a higher-intensity treatment for the high-risk group than for the lower-risk group.¹⁰

Panel B reports similar statistics for the NLSY97. In relative terms, the high-risk group remains disadvantaged in comparison to the lower-risk group. At the same time, the high-risk group had more-educated parents, higher family incomes, lower exposure

⁹All monetary amounts are expressed in thousands of 2018 dollars.

¹⁰The high- and lower-risk groups do not make up exactly 33 percent of the sample, as one might have expected from groups based on terciles. There are two reasons for this. First, in both generations, there is lumpiness in income/needs, owing to lumpiness in both reported household income and poverty thresholds. Second, the NLSY79 (i.e., absolute) terciles generally do not divide the NLSY97 sample into thirds.

to welfare, yet were more likely to live in a single-parent family, than the high-risk group in the NLSY79. It is conceivable that family incomes were higher, and welfare receipt was lower, due to the safety-net reforms themselves.

Predicted income/needs, our measure of risk, improved a bit between generations among the most disadvantaged group. In the earlier generation, mean predicted income/needs was 1.458 for women, and 1.490 for men, in the high-risk group. In the latter generation, that rose to 1.878 and 1.913, respectively. Predicted income/needs rose less between generations within the lower-risk group.

These differences in background characteristics are reflected in the differences across risk groups in the distribution of schooling. Figure 3.5 displays histograms of highest grade completed, disaggregated by sex, generation, and risk group.¹¹ For both women (Panel A) and men (Panel B), there is less mass to the right of the spike at 12 years among the high-risk group than among the lower-risk group. Also evident is the stark difference between these two groups and the lowest-risk group. Already in the earlier generation, the lowest-risk group was much more likely to complete more than twelve years of education, and particularly more likely to obtain at least 16 years, than the two higher risk groups. Combined with their families' higher income/needs during adolescence, their parents' higher education levels, and their minimal exposure to welfare in adolescence, this group seems like a poor comparison for the group that was most affected by safety-net reform. As a result, we omit them from the regression analysis below.

To provide a preview of the regression results, we present mean schooling outcomes by sex, generation, and risk group in Table 3.2, along with mean differences and differences-in-differences by group, restricting attention to the high-risk treatment group and the lower-risk comparison group.¹² The layout is similar to that of Table 3.1 above, with the

¹¹As in Table 3.1, these histograms are averaged over all sixteen sets of risk groups.

¹²Again, these tables are averaged over all sixteen sets of risk groups.

addition of rows and columns presenting the differences and differences-in-differences.

The first row of the Table shows that high-risk women in the earlier generation completed 11.947 years of education, leaving them 0.766 years behind lower-risk women, who completed 12.713 years. Both groups completed more schooling in the later generation, at 12.49 and 13.393 years, respectively. However, the gain among the high-risk group was less than that for the lower-risk group, such that the high-risk group fell behind by 0.137 years in relation to their lower-risk counterparts. This deficit is the result of decreasing rates of degree attainment, both at the bachelor's and sub-baccalaureate level.

High-risk men started out at an even greater disadvantage in relation to lower-risk men, but they gained more between generations. As a result, the high-risk men gained 0.161 years in relation to the lower-risk men. Contrary to women, men gained ground at all levels of degree attainment.

3.4.2 Outcome Regressions

The difference-in-difference regression model

We estimate the effect of safety-net reform via difference-in-difference outcome regressions that take the form:

$$y_{gi} = \alpha_0 + \alpha_1 A_g + \alpha_2 H_{gi} + \beta A_g H_{gi} + X_{gi} \gamma + u_{gi} \quad g = 1979, 1997; i = 1, \dots, n_g \quad (3.2)$$

where y_{gi} is a schooling outcome for the i th person in the g^{th} generation, such as highest grade completed; A_g is an “after” dummy, equal to one for members of the NLSY97 and equal to zero for members of the NLSY79; H_{gi} is a high-risk dummy, equal to one if the i th person in the g th generation is classified as high-risk, and equal to zero if she

is classified as lower-risk; and X_{gi} is a vector of controls. The terms α_0 , α_1 , α_2 , β , and γ are parameters to be estimated, and u_{gi} is an unobservable disturbance term. The coefficient β , on the interaction between the 1997-generation dummy and the high-risk dummy, is the coefficient of interest. Because we expect safety-net reform to affect males and females differently, we estimate equation (2) separately by sex.

In equation (2), α_2 captures pre-reform differences in schooling between the high- and lower-risk groups. The term α_1 captures general trends in educational attainment between the two generations. If the only difference between generations were safety-net reform, then β would identify the effect of safety-net reform as a bundle, meaning the joint effects of all programs, via both the resource and expectations channels.

Of course, there were other changes between generations that could have affected educational attainment. As discussed above, parental education rose and childhood living arrangements changed. We account for these changes by including in X_{gi} controls for maternal education, paternal education, family structure at age 12/14, family size, and race. We refer to this as our baseline specification.¹³

In addition, there were changes in a number of macro- and socio-economic factors between generations which may have independently influenced educational outcomes. These include real GDP/capita; unemployment; the return to education; inequality, as measured by the share of income accruing to the top 10 percent; and incarceration rates. In some specifications, we add these aggregate measures to the baseline specification; we refer to this as our baseline+macro specification. These measures reflect the values in the sample member's state of residence between the ages of 18 and 22. These controls help to better isolate the effect of safety-net reform.

For both the baseline and baseline+macro specifications, we further expand the sets

¹³The regressions also include missing value flags for these variables. These equal one when the corresponding variable is missing and zero otherwise. Missing values of the corresponding variable are recoded to zero.

of controls in two ways. First, we interact the control variables with the treatment-group dummy H_{gi} . These interactions allow the control variables to influence the outcome variables differently according to risk group. The combination of changing regressors across generations and differential effects by risk group could be falsely attributed to the effects of safety-net reform in the absence of these interactions. Second, we add interactions between all the regressors and the post-reform dummy A_g . These interactions account for otherwise unobserved trends in schooling that are associated with the effects of these factors which could otherwise bias our estimated effects of safety-net reform.

Inference

With sixteen sets of risk groups and six sets of controls variables, we estimate 96 different specifications of equation (2). We aggregate these estimates in two ways. First, we report the mean and median of the 96 estimated treatment effects. We then simulate the distributions of those test statistics using the resampling approach outlined by Simonsohn et al. (2020). We bootstrap the data 500 times, forcing the null hypothesis of no treatment effect and estimating all 96 specifications of the model at each bootstrap iteration. At each iteration, we calculate the mean and median bootstrap treatment effects.¹⁴ The 500 values so calculated provide estimates of the distributions of the mean and median treatment effects that would obtain under the null hypothesis of no effect. We calculate overall p-values for our test of the null hypothesis by counting the share of bootstrap means and medians that are at least as extreme as the actual mean and median.

Second, we plot the underlying estimates using the specification curve of Simonsohn et al. (2020). These graphs display the relationship between the 96 estimated treatment

¹⁴Bootstrapping is done at the level of the household, since some NLSY households contain multiple sample members.

effects and the characteristics of the regression models that produce them. We make one refinement to the scheme proposed by Simonsohn et al. (2020), plotting the magnitudes of the estimates, rather than their rank-ordering. This allows the reader to assess the magnitudes directly and also provides a sense of the variation among the estimates across specifications.

3.4.3 Main results

The first part of Table 3.3 summarizes estimates of β from equation (2). The first column presents results for women and the second presents results for men. In the first row, we present the mean of the estimates of the effect of safety-net reform on highest grade completed. The figure below that is the simulated p-value for the mean, obtained as described above. The next rows present the median estimate and its simulated p-value.¹⁵ In addition to the variables shown, the underlying regressions included various sets of control variables as described above.

The mean estimate for women is -0.175, with a p-value of 0.252, and the median estimate is -0.221, with a p-value of 0.010. Twenty-five of the 96 underlying coefficients were significant at the 5 percent level.¹⁶ If the estimates were independent, the probability that 25 of 96 coefficients would be significant would be 6.9×10^{-12} under the null of no effect. The difference between the binomial and simulated p-values reflects the substantial dependence among the estimates. Overall, these results suggest that safety-net reform may have modestly reduced completed schooling for women, although that reduction is not statistically significant.

For men, the central tendency of the estimated treatment effects on years of education

¹⁵All analyses were carried out using Stata 16. For an even number of observations k , Stata defines the median as the average of the $k/2$ th and $(k+2)/2$ th values.

¹⁶Both here and below, significance levels for individual regression coefficients are based on standard errors that are clustered at the level of the household.

is positive. The mean estimate is 0.171 [p-value=0.046] and the median is 0.134 [0.074]. Safety-net reform increased highest grade completed among men by a modest, marginally significant amount.

Figure 3.1 plots the underlying estimated treatment effects using specification curves Simonsohn et al. 2020. Each point in the curve plots an estimate of β from equation (2). Among all the rows representing a particular specification characteristic, there are 96 estimates; the number of points on each row is equal to $96/p$, where p is the number of values taken by that characteristic. Thus there are 48 estimates plotted on each of the two rows for the choice of dependent variable in the risk regression (Average (income/needs) vs. (Average income)/(Average needs)), and 16 estimates plotted on each of the six rows pertaining to the specification of the outcome regression. Panel A presents the curve for women and Panel B presents the curve for men.

Figure 3.1: Estimates of the effect of safety-net reform on highest grade completed, by sex and specification characteristics



Notes: Each point plots an estimated treatment effect (i.e., an estimate of β from equation (2)). Among all the rows representing a particular specification characteristic, there are 96 estimates; the number of points on each row is equal to $96/p$, where p is the number of values taken by that characteristic. Thus there are 48 estimates plotted on each of the two rows for the choice of dependent variable in the risk regression (Average (income/needs) vs. (Average income)/(average needs)), and 16 estimates plotted on each of the six rows pertaining to the specification of the outcome regression. The solid vertical line depicts the mean estimate, and the dashed vertical line depicts the median estimate.

As a general matter, the specification curves convey several pieces of information. First, the mean treatment effect of -0.175, presented in the first row of Table 3.3, is depicted with a solid vertical line, and the median treatment effect of -0.221, presented in the third row of Table 3.3, is marked with a dashed vertical line. Second, by plotting the magnitudes of the coefficients, rather than their rank ordering, as in Simonsohn et al. (2020), the specification curves depict the range of variation of the estimates. For women, the estimated treatment effects range from roughly -0.6 to 0.6, with a majority being negative.

Third, they show which specification characteristics have the greatest influence on the estimates. To see this, note that all 96 estimates are plotted on each set of rows representing a single specification characteristic. If the estimates are insensitive to the different values of the characteristic, then the estimates plotted on the associated rows will overlap. Less overlap indicates greater sensitivity.

The first two rows of Figure 3.1 show estimates according to the dependent variable used to estimate the risk regression, and the second two rows show them according to whether the risk regression was weighted or not. In both cases, the estimates overlap to a great extent. Put differently, basing the risk regression on mean (income/needs) vs. (mean income)/(mean needs) has little effect on the estimates. The same can be said for the choice of estimator, and likewise, in the next two rows, for the use of relative vs. absolute thresholds to define the risk groups. The next two rows show that the largest positive estimates tend to come from specifications that include family structure.

Which variables are included in the outcome regressions is more consequential. The specifications with no interactions, either by risk-group or by generation, tend to yield the most negative estimates. Those with generation interactions tend to be smaller in absolute value, but most remain negative.

The specification that stands out most includes the baseline+macro variables plus

risk-group interactions. This single specification delivers most of the positive estimates in the Figure. This shows that it is difficult to disentangle the effects of safety-net reform from those of other macro- and socio-economic factors whose effects vary by risk group.

The specification curve for men, in panel B, differs in several regards. First, the estimates span a smaller range, extending from roughly -0.05 to 0.20.¹⁷ Second, a much larger share of the estimates is positive. Finally, in contrast to the estimates for women, the specification that yields the largest positive estimates now involves a combination of baseline+macro regressors, risk-group interactions, and the *exclusion* of family structure variables.

The remainder of Table 3.3 presents estimates of the effect of safety-net reform on three additional educational outcomes. These are high school graduation, post-secondary sub-baccalaureate degree completion, and bachelor's degree attainment. The estimates are all based on linear probability models. These measures of degree attainment were collected from a different set of questions than those used to measure highest grade completed.

The estimates in Table 3.3 show that safety-net reform had little effect on high school graduation. The mean and median estimates for women are both negative but fairly small, and neither is significant. The estimates for men are both positive and small, and likewise insignificant.


Figure 3.2 presents the relevant specification curves. The range of estimates for women is fairly narrow, with the great majority of estimates being negative. The combination of baseline+macro explanatory variables, plus risk group interactions, plus the exclusion of family structure, is associated with the largest negative estimates. Estimates for men span a broader range and are more evenly distributed between positive and negative

¹⁷Here and in many cases to follow, we observe that the group whose estimates have a smaller range also has smaller simulated p-values, even though the number of significant coefficients is similar.


numbers. For men, the combination of baseline+macro explanatory variables, risk group interactions, and the exclusion of family structure leads to the largest positive estimates.

Figure 3.2: Estimates of the effect of safety-net reform on high school graduation, by sex and specification characteristics

(a) Women



(b) Men



Note: Each point plots an estimated treatment effect (i.e., an estimate of β from equation (2)). Among all the rows representing a particular specification characteristic, there are 96 estimates; the number of points on each row is equal to the $96/p$, where p is the number of values taken by that characteristic. Thus there are 48 estimates plotted on each of the two rows for the choice of dependent variable in the risk regression (Average (income/needs) vs. (Average income)/(average needs)), and 16 estimates plotted on each of the six rows pertaining to the specification of the outcome regression. The solid vertical line depicts the mean estimate, and the dashed vertical line depicts the median estimate.

The next set of rows in Table 3.3 show the estimated effect of safety-net reform on the attainment of a post-secondary credential below the level of a bachelor's degree, such as an associate's degree. For women, the estimated effects are small and negative. The median estimate [p-value] is -0.035 [0.128] and the mean estimate is -0.035 [0.070]. For men, the estimated effects are small and positive. The mean effect is 0.026 and the median is 0.027, with p-values of 0.034 and 0.032, respectively.

Figure 3.3 presents the associated specification curves. Here, the range of estimates is somewhat broader for women than for men. Most of the estimates for women are negative, whereas most of the estimates for men are positive. For women, the combination of baseline+macro explanatory variables, risk group interactions, and the exclusion of family structure now leads to the largest positive estimates. That said, fewer of the estimates in this group appear to be outliers than was the case for either highest grade completed or high school graduation. For men, most of the estimates are positive. The estimates from specifications involving risk-group interactions stand out here less than they did for the other attainment measures. Here, the specification characteristic that leads to the greatest variation in estimates involves the inclusion/exclusion of family background.

Figure 3.3: : Estimates of the effect of safety-net reform on Post-secondary, sub-baccalaureate credential, by sex and specification characteristics




Note: The graph shows how the five specification characteristics (four binary and one six-valued) influence the resulting 96 estimated treatment effects. Each point plots an estimate of β from equation (2). Among all the rows representing a particular specification characteristic, there are 96 estimates; the number of points on each row is equal to $96/p$, where p is the number of values taken by that characteristic. Thus there are 48 estimates plotted on each of the two rows for the binary specification characteristics (such as the dependent variable used in the risk regression), and 16 estimates plotted on each of the six rows pertaining to the specification of the outcome regression. The solid vertical line depicts the mean estimate, and the dashed vertical line depicts the median estimate.

The last set of rows in Table 3.3 show the estimated effect of safety-net reform on bachelor's degree completion. For women, the mean and median estimates are negative, fairly small, and insignificant. For men, they are positive, smaller, and marginally significant.


Figure 3.4 displays the relevant specification curves. The range of estimates is similar for men and women. A majority of estimates for women is negative, whereas the majority for men is positive. For women, specifications without interactions now tend to yield larger negative estimates, as do specifications that include family structure. For men, the role of both interactions and family structure is smaller than was the case for the other outcomes.

Figure 3.4: : Estimates of the effect of safety-net reform on BA plus, by sex and specification characteristics

(a) Women



(b) Men



Note: Each point plots an estimated treatment effect (i.e., an estimate of β from equation (2)). Among all the rows representing a particular specification characteristic, there are 96 estimates; the number of points on each row is equal to the $96/p$, where p is the number of values taken by that characteristic. Thus there are 48 estimates plotted on each of the two rows for the choice of dependent variable in the risk regression (Average (income/needs) vs. (Average income)/(average needs)), and 16 estimates plotted on each of the six rows pertaining to the specification of the outcome regression. The solid vertical line depicts the mean estimate, and the dashed vertical line depicts the median estimate.

In terms of substance, the summary results for completed schooling for women are all negative, with the reductions coming from decreases in both high school graduation and post-secondary degree completion. That said, it can be difficult to separate the effects of safety-net reform from those of trending control variables whose coefficients vary by risk group. Moreover, all but a few of the summary estimates presented in Table 3.3 are insignificant, even at the 10 percent level. For men, the summary estimates are all positive. The estimates pertaining to high school graduation are insignificant, but those involving higher degrees are mostly significant.

Methodologically, the specification curves show that certain characteristics are associated with variation in the estimated treatment effects, but that that variation is not uniform. In the case of high school graduation, including family structure, particularly in conjunction with risk-group interactions involving the macro- and socio-economic variables, resulted in the largest positive estimates. However, for men, excluding family structure, in conjunction with the same other variables, resulted the largest positive estimates. Characteristics which led to pronounced variability for highest grade completed and high school graduation were associated with less variability for bachelor's degrees.

We see this as a virtue of using the specification curve approach, rather than attempting to find a single, “best,” specification. It is clear that any single specification could result in a set of estimates that differed greatly from those that we report here. At the same time, those estimates would not necessarily differ all in one direction, being more or less positive, for example. In isolation, one can make a case for any of the specification choices that we have made. Assessing the validity of a single combination of several choices is a much harder task, leading to considerable model uncertainty. The specification curve approach helps provide a transparent means of accounting for that uncertainty in drawing inference.

3.4.4 Additional outcomes: labor supply and post-secondary enrollment

Tables 3.4 and 3.5 provide some additional insights into what high-risk youths were doing between the ages of 18 and 22, the time when college-goers tend to be enrolled in school. We first consider their labor supply, then analyze various combinations of employment and enrollment. For these analyses, we estimate regressions similar to equation (2), except now there are multiple observations per person, with one observation per person per year. Accordingly, we included the annual measure of the macro- and socio-economic variables associated with each year's observation, rather than averaging them over the ages 18-22. We report only summary tables of regression results here, foregoing the specification curves in order to save space.

Table 3.4 presents estimates for three dependent variables: an indicator for any employment during the year, the number of weeks worked during the year, and the number of hours worked during the year. Safety-net reform increased women's labor supply over this age range, as shown in Panel A. They were roughly 6 percentage points more likely to be employed; 50 of the 96 underlying estimates were significant, and p-values for both the mean and median coefficients equaled zero to three decimal places. They worked roughly two more weeks a year, which translated to 70 to 86 hours/year. All of these estimates have p-values less than 0.05.

For men, the story was quite different. Safety net reform barely changed their labor supply, however measured. Moreover, almost none of the estimated effects were significant.

The results in Table 3.5 reinforce these findings. Here we consider various combinations of enrollment and employment status as of May 1 in each year. These measures allow us to analyze substitution between work and schooling as a result of safety-net

reform.

For 18- to 22-year-old women, safety-net reform led to a shift away from enrollment and toward employment. It reduced the likelihood of working while enrolled by roughly 4 percentage points, while increasing the probability of working while not enrolled by approximately 10 percentage points. It also reduced the likelihood of neither working nor being enrolled, which is sometimes viewed as a measure of economic idleness among young adults, by 7.3 percentage points. All of these results have p-values less than 5 percent. Safety-net reform had essentially no effect on the likelihood of not working while enrolled.

For young men, in contrast, safety-net reform had little effect on most combinations of work and enrollment. All of the estimates are small, and most are insignificant. The exception involves economic idleness, which fell by about 2.5 percentage points.

In Table 3.6, we interact employment and enrollment with the presence of children. The employment interactions provide a type of falsification test. Both welfare reform and the EITC provide work incentives primarily for people with children. If the various employment effects we estimate are attributable to safety-net reform, they should be concentrated among workers with children. If we found employment effects among people without children, it would cast doubt on whether they truly stemmed from changes in the safety net.

In fact, safety-net reform raised the likelihood of working and having children for women by about 6 percentage points; both the mean and median estimates are significant. It had essentially no effect on the employment of young women without children. It also had small but significant positive effects on the employment of young men with children.

How the effect of safety-net reform on enrollment should vary between young adults with and without children is not entirely clear. As an empirical matter, such effects as there are are negative and concentrated among young women without kids. However,

those estimates are at best marginally significant.

These early effects on educational attainment and labor supply raise questions about other outcomes and later points in the life cycle. In Bastian et al. (forthcoming), we build on the current paper, analyzing employment and safety-net participation as well as marriage and childbearing. For women, we find null effects on receiving cash welfare, but positive effects on receiving SNAP and EITC benefits. For men, we find largely null effects.

3.4.5 Other dimensions of heterogeneity

One might expect the effects of safety-net reform to vary along a number of observable dimensions, such as family structure during adolescence, family size, and birthdate. Sample members who grew up with single parents would have been more likely to receive welfare or EITC payments; both welfare benefits and EITC payments are a function of family size; and several of the most important Medicaid expansions were reserved for children born after September 30, 1983. In an attempt to estimate treatment effect heterogeneity, we divided the sample into subgroups along each of these dimensions. For family structure we compare sample members from one- and two-parent families. For birthdate we divided the NLSY97 into those born before or after September 30, 1983. For family size, we split the sample at four individuals, which divided it most nearly in two.

The problem is that the resulting sample sizes were fairly small. As a result, of the total of 8640 between-group differences estimated (15 outcomes x 2 sexes x 96 specifications x 3 sets of subgroups), only 425, or 4.9 percent, of the between-group differences were significant at the five percent level. Restricting attention to women, only 126, or 2.9 percent, of the 4320 between-group differences were significant. The subgroup estimates

seem to be dominated by noise. As a result, we do not report them.

3.5 Discussion

To summarize our results, we find evidence that safety-net reform may have reduced educational attainment at age 25 among women who were adolescents at the time the reforms were enacted. Completed schooling fell by about one-fifth of a year, and high school graduation, sub-baccalaureate degree completion, and college graduation all fell by about 3 percentage points. Enrollment between the ages of 18 and 22 generally fell as well. That said, although all the evidence points in the same direction, most of it is insignificant. It is particularly difficult to distinguish the effects of safety-net reform from those of various macroeconomic factors whose effects vary by risk group.

What's clearer is that safety-net reform increased young women's work. Employment rose by about 6 percentage points, with substantial gains in both weeks and hours of work. Working rose most among those not enrolled, whereas working while enrolled fell. Economic idleness, defined as neither working nor being enrolled, also fell. Furthermore, most of the increase in employment occurred among young women with children, who stood to gain most from the EITC and lose the most from welfare reform. Nearly all of these estimates were significant at least at the 5 percent level.


The picture was much different for young men. Years of schooling rose by about one-seventh of a year, with small gains in both sub-baccalaureate degrees and college completion. Most of these estimates were significant at the 5 percent level.

At the same time, safety-net reform had little effect on work among young men. It had no significant effect on their employment, hours, or weeks of work. However, it did reduce economic idleness, and increase employment among those with children, by small amounts.


Our goal in the rest of this section is to relate these findings to prior work and offer an explanation for them. We first note that the increases in highest grade completed shown in Figure 3.5 and Table 3.2 are consistent with research showing that aggregate levels of education increased between our two generations Murnane 2013. Most children in both the treatment and comparison groups grew up in the bottom half of the income distribution, and may have benefited from favorable school funding trends, consistent with Jackson et al. (2016) and Lafortune et al. (2018).

Figure 3.5: Distribution of highest grade completed, by sex, generation, and risk group

(a) Women



(b) Men



Notes: Averaged over all 16 risk groups.

We next ask whether our results merely represent the extension into early adulthood

of the negative consequences of welfare reform that arose while our post-reform generation was still in its teens. As discussed above, some early studies of the effects of welfare reform indicated that children who were teens at the time suffered as a result, particularly in terms of scholastic outcomes Shaefer et al. 2020. One important problem with this explanation is that previous evidence has shown either that welfare reform adversely affected adolescents similarly by sex, or that it had worse consequences for boys than for girls (Gennetian et al. (2002); Dave et al 2020). In contrast, we find reductions in schooling for young women, but not for young men. Furthermore, it is not clear whether one should consider sizeable increases in employment among young women, and reductions in economic idleness, as adverse effects.

Thus the main question is how we reconcile our results with several studies that have recently found the EITC and Medicaid expansions to have small but favorable effects on the educational attainment of children whose families were affected by those expansions. As discussed above, those studies focus implicitly on the resource channel, that is, on the extent to which safety-net reform affected children's educational attainment by raising their parents' income while the children lived at home. We have no reason to dispute those findings, and indeed our results for males are consistent with them. Thus what we seek to explain is how safety-net reform could reduce young women's schooling, even though the component of its effect that worked solely through the resource channel was positive.

A key difference between our approach and that of previous studies is that our estimates should reflect both resource and expectations channels. Thus one reconciliation between our results and those of others involves expectations. Above we argued that the expectations channel should be more important to females than males, because women historically were much more likely to receive welfare than men, and roughly 60 percent of current adult EITC beneficiaries are women Hoynes & Patel 2018. Safety-net reform

increased young men's schooling, consistent with positive resource effects, and otherwise had little effect on their behavior. For women, the expectations channel may have resulted in a negative effect on schooling that dominated the positive effect that stemmed from the resource channel.

The next question is, why might the expectations channel have a negative effect on educational attainment, when safety-net reform should have raised the return to human capital? Here we offer several admittedly tentative explanations. First, whether the increase in the returns to human capital should take the form of schooling or labor market experience is not clear *ex ante*, but depends on the return to each type of human capital. Second, at the same time that safety-net reform increased the return to schooling, the wage subsidy implicit in the EITC also raised the return to work. If schooling and work are rivalrous activities, as in Heckman et al. (2002), the result may have been more work rather than more schooling. Finally, the decision to work may also have been promoted by the message conveyed during the welfare reform era, which generally favored work over education Haskins 2006, Moffitt 2007.

Although these explanations generally fit the data, we acknowledge that direct tests would be more convincing. At any rate, more work on the topic is essential. Findings that safety-net reform has had positive effects that arise via the resource channel are not sufficient to conclude that the long-term effects of safety-net reform have been favorable. Reductions in schooling that stem from the expectations channel could eventually have adverse effects on earnings, despite their favorable short-term effects on early-career employment. We build on our results here in Bastian et al. (forthcoming), where we look at employment, safety-net participation, and family formation through age 32. Future research should consider additional outcomes and even longer-run effects.

Table 3.1: Explanatory variable means by generation, sex and risk group

Panel A: NLSY 79									
Variable	Women				Men				Total
	High-risk (Treatment) (1)	Lower-risk (Comparison) (2)	Lowest-risk (Not used) (3)	Total (4)	High-risk (Treatment) (1)	Lower-risk (Comparison) (2)	Lowest-risk (Not used) (3)	Total (4)	
Mom's highest grade completed	9.588	11.487	13.417	11.467	9.675	11.569	13.471	11.571	
Dad's highest grade completed	8.476	11.196	14.345	11.628	8.724	11.294	14.468	11.788	
Family size	6.165	4.855	4.108	5.063	6.267	4.881	4.234	5.141	
HH income, wave 0	38.037	59.945	78.149	57.927	37.922	56.187	78.293	56.750	
On welfare, wave 0	0.320	0.081	0.022	0.142	0.304	0.094	0.023	0.142	
Black	0.650	0.303	0.197	0.386	0.636	0.330	0.175	0.384	
Lived with both parents at 12/14	0.420	0.718	0.815	0.648	0.443	0.730	0.822	0.663	
Risk index	1.458	2.518	3.731	2.537	1.490	2.538	3.743	2.565	
N	667	711	615	1,993	705	743	662	2,110	
Panel B: NLSY 97									
Variable	Women				Men				Total
	High-risk (Treatment) (1)	Lower-risk (Comparison) (2)	Lowest-risk (Not used) (3)	Total (4)	High-risk (Treatment) (1)	Lower-risk (Comparison) (2)	Lowest-risk (Not used) (3)	Total (4)	
Mom's highest grade completed	10.987	12.382	14.682	12.977	11.072	12.461	14.632	13.020	
Dad's highest grade completed	10.814	12.035	14.793	13.125	10.711	12.027	14.801	13.142	
Family size	5.381	4.202	4.043	4.431	5.185	4.289	4.088	4.434	
HH income, wave 0	44.990	57.451	95.595	70.372	45.503	61.303	95.046	71.913	
On welfare, wave 0	0.272	0.101	0.027	0.111	0.232	0.091	0.019	0.096	
Black	0.524	0.374	0.186	0.337	0.514	0.358	0.175	0.323	
Lived with both parents at 12/14	0.159	0.327	0.670	0.421	0.161	0.375	0.705	0.456	
Risk index	1.878	2.668	4.039	3.005	1.913	2.704	4.049	3.046	
N	855	1,241	1,360	3,455	914	1,223	1,474	3,611	

Notes: Tables are averaged over all 16 sets of risk groups. *N* indicates the mean sample size.

Table 3.2: Outcome variable means by generation, sex and risk group

		Treatment		Comparison		Difference		Treatment		Comparison		Difference	
NLSY79, women								NLSY79, men					
Highest grade completed	11.947			12.713				11.473			12.474		
High school graduation	0.813			0.911				0.713			0.867		
Post-secondary, sub-baccalaureate credential	0.121			0.238				0.086			0.194		
BA or greater	0.054			0.142				0.044			0.134		
<i>N</i>	667			710				705			742		
NLSY97, women								NLSY97, men					
Highest grade completed	12.490			13.393				11.838			12.678		
High school graduation	0.817			0.907				0.773			0.884		
Post-secondary, sub-baccalaureate credential	0.165			0.297				0.114			0.187		
BA or greater	0.108			0.217				0.070			0.132		
<i>N</i>	791			1,134				843			1,107		
Difference													
Highest grade completed	0.543			0.679				0.365			0.203		
High school graduation	0.005			-0.004				0.060			0.017		
Post-secondary, sub-baccalaureate credential	0.044			0.059				0.028			-0.007		
BA or greater	0.054			0.075				0.025			-0.003		

Notes: Tables are averaged over all 16 sets of risk groups. *N* indicates the mean sample size for highest grade completed. Sample sizes for the educational attainment indicators (not shown) are slightly smaller because they are derived from different questionnaire items with slightly greater non-response.

Table 3.3: Estimated effects of safety-net reform on educational outcomes

	Women (1)	Men (2)
Highest grade completed		
Mean estimate	-0.175	0.171
P-value	0.252	0.046
Median estimate	-0.221	0.134
P-value	0.100	0.074
No. of signif. coeffs. (out of 96)	25	25
High school graduation		
Mean estimate	-0.031	0.018
P-value	0.112	0.252
Median estimate	-0.028	0.003
P-value	0.120	0.500
No. of signif. coeffs. (out of 96)	0	6
Post-secondary, sub-baccalaureate credential		
Mean estimate	-0.035	0.026
P-value	0.128	0.034
Median estimate	-0.039	0.027
P-value	0.070	0.032
No. of signif. coeffs. (out of 96)	21	7
BA or greater		
Mean estimate	-0.032	0.021
P-value	0.136	0.044
Median estimate	-0.034	0.016
P-value	0.076	0.106
No. of signif. coeffs. (out of 96)	21	8

Notes: Based on 96 regressions per outcome. Sample sizes range from 2,598 to 3,601 for women and from 2,683 to 3,690 for men. Differences in sample sizes across regressions stem from differences in the definitions and sizes of risk groups, and from differences in the number of missing values across outcomes.

Table 3.4: Estimated effects of safety-net reform on labor supply, ages 18–22

	Women (1)	Men (2)
Any weeks worked		
Mean estimate	0.057	0.001
P-value	0.000	0.470
Median estimate	0.064	-0.002
P-value	0.000	0.452
No. of signif. coeffs. (out of 96)	50	2
Weeks worked		
Mean estimate	2.089	-0.294
P-value	0.016	0.456
Median estimate	1.883	-0.390
P-value	0.026	0.376
No. of signif. coeffs. (out of 96)	37	6
Hours worked		
Mean estimate	86.437	-19.111
P-value	0.016	0.404
Median estimate	70.395	-30.498
P-value	0.042	0.298
No. of signif. coeffs. (out of 96)	42	12

Notes: Based on 96 regressions per outcome. Sample sizes range from 12,478 to 17,999 for women and from 12,082 to 18,363 for men. Differences in sample sizes across regressions stem from differences in the definitions and sizes of risk groups, and from differences in the number of missing values across outcomes.

Table 3.5: Estimated effects of safety-net reform on economic activity, ages 18–22

	Women (1)	Men (2)
Working, enrolled		
Mean estimate	-0.044	-0.007
P-value	0.018	0.494
Median estimate	-0.039	-0.008
P-value	0.010	0.440
No. of signif. coeffs. (out of 96)	34	23
Working, not enrolled		
Mean estimate	0.111	0.022
P-value	0.000	0.182
Median estimate	0.098	0.026
P-value	0.000	0.130
No. of signif. coeffs. (out of 96)	66	14
Not working, enrolled		
Mean estimate	0.006	0.011
P-value	0.210	0.150
Median estimate	0.006	0.003
P-value	0.210	0.418
No. of signif. coeffs. (out of 96)	0	17
Not working, not enrolled		
Mean estimate	-0.073	-0.025
P-value	0.000	0.044
Median estimate	-0.073	-0.027
P-value	0.000	0.034
No. of signif. coeffs. (out of 96)	57	10

Notes: Based on 96 regressions per outcome. Sample sizes range from 12,717 to 17,998 for women and from 12,430 to 18,407 for men. Differences in sample sizes across regressions stem from differences in the definitions and sizes of risk groups, and from differences in the number of missing values across outcomes.

Table 3.6: Estimated effects of safety-net reform on economic activity by presence of children, ages 18–22

	Women (1)	Men (2)
Working, with children		
Mean estimate	0.064	0.028
P-value	0.000	0.012
Median estimate	0.067	0.032
P-value	0.000	0.002
No. of signif. coeffs. (out of 96)	50	32
Working, without children		
Mean estimate	0.002	-0.013
P-value	0.410	0.290
Median estimate	-0.006	-0.014
P-value	0.436	0.258
No. of signif. coeffs. (out of 96)	5	2
Enrolled, with children		
Mean estimate	-0.002	-0.002
P-value	0.410	0.242
Median estimate	-0.002	-0.003
P-value	0.360	0.168
No. of signif. coeffs. (out of 96)	0	2
Enrolled, without children		
Mean estimate	-0.040	0.005
P-value	0.064	0.250
Median estimate	-0.033	-0.011
P-value	0.078	0.366
No. of signif. coeffs. (out of 96)	24	22

Notes: Based on 96 regressions per outcome. Sample sizes range from 12,480 to 17,688 for women and from 12,282 to 18,243 for men. Differences in sample sizes across regressions stem from differences in the definitions and sizes of risk groups, and from differences in the number of missing values across outcomes.

Appendix A

Appendix for Chapter 1

A.0.1 Conceptual Model on Differential Health Impacts

This subsection begins with a conceptual framework that explains how heat affects health through adaptation and/or the biological channel. Next, I introduce a set of empirical regressions that test for potential differential effects by heat composition. Combining these two parts, adaptation provides a more plausible explanation for the differential health impacts compared to the biological channel.

Suppose urban heat UH and background temperature T impose the same health risks. Then overall temperature L , the sum of the two, would be sufficient for evaluating the health-temperature relationship, and introducing a composition variable should have minimal impact unless health risks indeed differ by heat composition.

Following Grossman (1972), Becker & Murphy (1988), Moretti & Neidell (2011), Barreca et al. (2016), and Deschênes et al. (2017), I formulate a health production model that accounts for temperatures, other factors (W), and avoidance behavior (A):

$$H = H[UH, T, W, A(UH, T)] \tag{A.1}$$

For $x \in \{UH, T\}$, the total differentiation is as follows:

$$\frac{dH}{dx} = \underbrace{\frac{\partial H}{\partial x}}_{B_x} + \underbrace{\frac{\partial H}{\partial A} \frac{\partial A}{\partial x}}_{A_x} \quad (\text{A.2})$$

The overall impact of x on health consists of two components: the direct effect through the biological channel (B_x) and the indirect effect through the avoidance channel (A_x).

In the existing literature, overall temperature (L) has been considered the primary determinant of health impacts, without accounting for its composition. In other words, potential differences between urban heat (UH) and background temperature (T) in the health-temperature relationship have not been examined. Without loss of generality, consider the simplest regression:

$$H = L\beta_0 + W'\gamma_0 + e_0 \quad (\text{A.3})$$

where W represents a set of controls, including fixed effects. Next, I consider the following regression, which explicitly accounts for UH and is mathematically equivalent to including both UH and T on the right-hand side, as shown in Proposition 2. I rely on A.3 because it aligns with existing approaches and provides a more straightforward interpretation.

$$H = L\beta_1 + UHr_1 + W'\gamma_1 + e_1 \quad (\text{A.4})$$

Using equations A.3 and A.4, the following proposition evaluates whether urban heat needs to be explicitly accounted for.

Proposition 2 $\beta_0 \neq \beta_1$ implies $B_{UH} + A_{UH} \neq B_T + A_T$. In other words, changes in the estimate of L after including UH controls suggest differential total effects of UH and T, summing across biological and avoidance channels.

Proof. By the Frisch–Waugh–Lovell Theorem, β_1 estimates the same relationship between H and L after both have been adjusted for the influence of UH , that is, $\frac{dH}{dL}|\overline{UH}$.

When $\beta_1 \neq \beta_0$, implying $\frac{dH}{dL} \neq \frac{dH}{dL}|\overline{UH}$, we have:

$$\begin{aligned}\frac{dH}{dL} &= \frac{\partial H}{\partial T} \frac{\partial T}{\partial L} + \frac{\partial H}{\partial UH} \frac{\partial T}{\partial UH} + \frac{\partial H}{\partial A} \left(\frac{\partial A}{\partial T} \frac{\partial T}{\partial L} + \frac{\partial A}{\partial UH} \frac{\partial UH}{\partial L} \right) \\ &= \alpha \frac{dH}{dT} + (1 - \alpha) \frac{dH}{dUH}\end{aligned}\tag{A.5}$$

with $\alpha := \frac{\partial T}{\partial L}$. Since $1 = \frac{\partial L}{\partial L} = \frac{\partial(UH+T)}{\partial L}$, we obtain $\frac{\partial UH}{\partial L} = 1 - \alpha$. Similarly:

$$\frac{dH}{dL}|\overline{UH} = \frac{\partial H}{\partial T} \frac{\partial T}{\partial L} + \frac{\partial H}{\partial A} \frac{\partial A}{\partial T} \frac{\partial T}{\partial L} = \frac{dH}{dT}\tag{A.6}$$

This implies that:

$$\beta_1 \neq \beta_0 \Rightarrow \frac{dH}{dUH} \neq \frac{dH}{dT} \Rightarrow B_{UH} + A_{UH} \neq B_T + A_T$$

Q.E.D.

Equation A.5 shows that the effects of L on health are a weighted average of the effects of UH and T , with weights determined by changes in heat composition (α). That is, composition matters unless the marginal effects of the two are identical. Furthermore, when the regression controls for UH , Equation A.6 indicates that the estimate of L now represents only the effects of T . Since UH is held constant, all effects of L are attributed solely to T , capturing its influence exclusively.

As discussed in Section 1.2, urban heat UH and background temperature T differ in many aspects, including geographical scale, formation, and, most importantly, adaptation potential. Given the complexity of urban climate dynamics, the same adaptation method may not be equally effective for heat events that are homogeneous across neighborhoods

versus those that are highly localized. Thus, it is crucial to examine the potential differential impacts of these two heat components and assess whether explicitly incorporating urban heat improves modeling of the health-temperature relationship in urban settings.

A.0.2 Urban Heat Mitigation Simulation

I begin by simulating the health benefits associated with different levels of urban heat using urban heat estimates. Next, I compare these results to simulations based on overall temperature estimates to demonstrate that heat composition matters in policy evaluation. Finally, I examine tree planting, a widely used mitigation policy implemented in numerous U.S. cities.

The first exercise explores the extent to which urban heat mitigation improves health outcomes. For this analysis, I rely on UH estimates from Equation 1.4. The benchmark is the predicted caseload $\hat{y}_b$ under actual urban heat levels, while the comparisons involve four hypothetical urban heat reduction scenarios, $\hat{y}_{h1}$ to $\hat{y}_{h4}$. I calculate the predicted proportional change in caseload ($\frac{\hat{y}_{h4}-\hat{y}_b}{\hat{y}_b}$) for all observations in my sample. Then, I aggregate them using 1-degree background temperature bins and plot the results, as shown in Figure 1.9.

The second exercise evaluates whether accounting for heat composition affects policy outcomes. I apply the same mitigation scenarios but modify the simulation model to one based on overall temperature, i.e., Equation 1.2. Using the same procedure, I compute proportional changes in predicted caseload at each 1-degree background temperature bin and compare these reductions to those from the first exercise in Figure 1.10.

Finally, I link these results to policy variables by incorporating tree canopy data from the U.S. Forest Service, Department of Agriculture. To maintain consistency with the previous simulations, I estimate the effects of trees on urban heat using a separate model


with both cross-validated random forest and neural network methods, both of which produce similar results. Specifically, this estimation uses panel data on urban heat and background temperature from 2018, along with cross-sectional tree canopy data from the same year. The model captures the effect of static tree canopy on urban heat at the block-group level, allowing the impact to vary by season, location, and background temperature. Calendar-year effects are not included, as tree canopy data are cross-sectional. However, Table 1.1 indicates that annual trends have minimal impact on urban heat.

For the simulation, I vary tree canopy levels to estimate the corresponding urban heat reductions, predict the caseloads using these adjusted urban heat levels as $\hat{y}_h$, and then follow the same procedure to compute proportional changes in caseload. Two tree planting strategies are considered: the first uniformly increases tree canopy by a fixed proportion, while the second targets the lower end of the distribution, raising tree canopy to specific thresholds. For each strategy, I vary the amount of additional tree cover and report the resulting proportional changes in predicted caseload in Figure 1.11.

Health impacts are estimated using two approaches: (1) group-mean estimates (i.e., indicator variables for each UH group, resulting in a constant risk level within each UH bin) and (2) a more flexible specification that allows group-specific slopes. I rely on the second specification for reporting results, as the first approach involves greater information loss, which varies systematically across mitigation scenarios. To see this, consider a conservative UH mitigation strategy such that urban heat reductions are small, e.g., from 2.99 to 1.01. In such cases, the predicted health impact is minimal because both values fall within the same UH bin. This issue is less pronounced with more aggressive mitigation strategies but still results in an underestimation of the true effects. To ensure comparability and accuracy, I present results from the second specification, which, while similar to the first, offers a more precise estimate—especially for conservative UH mitigation approaches.


A.0.3 Appendix Figures and Tables

Figure A.1: Sample Areas



Notes: This figure overlays urban areas and urban buffers used in the urban heat calculation with the sample areas used in the analysis. These areas differ because urban heat calculations are based solely on land type, whereas the analysis sample is defined as the intersection of 2010 Census Urban Areas and EMS service areas.


Figure A.2: Robustness Check: Contemporaneous Effects of L in Alternative Specifications




Notes: Robustness check for Figure 1.3. Within each panel, the colored line replicates the main estimates from the text, while the two gray lines report estimates with block-group-month and/or block-group-week fixed effects (instead of block-group fixed effects). The last panel presents estimates using the number of EMS dispatches as the outcome variable, applying the same set of fixed effects.

Figure A.3: Robustness Check: Contemporaneous Effects of L by High-Low UH in Alternative Specifications

(a) Cross-Block-Group Comparison




(b) Within-Block-Group Comparison




Notes: Robustness check for Figure 1.5. In each panel, the top-left graph reproduces the main figure from the text, while the next two incorporate block-group-month and/or block-group-week fixed effects. The final panel presents estimates using the number of EMS dispatches as the outcome variable with the same set of fixed effects.

Figure A.4: Robustness Check: Contemporaneous Effects of Urban Heat in Alternative Specifications




Notes: Robustness check for Figure 1.6. Within each panel, the colored line replicates the main estimates, while the two gray lines report estimates incorporating block-group-month and/or block-group-week fixed effects. The final panel presents estimates using the number of EMS dispatches as the outcome variable with the same set of fixed effects.

Figure A.5: Robustness Check: Effects of Urban Heat by T 


Notes: Estimates are based on Equation 1.4, with treatments varying by T . The model controls for city-by-date, calendar-year, and block-group fixed effects. City-by-date fixed effects absorb T and other regional macroclimatic conditions that vary at the same level. Background temperatures around 70 degrees roughly separate the sample into the bottom 40% and the top 60%. Confidence intervals represent 95% standard errors, clustered at the census block level. The reference group is $UH \in [-1, 1]$ (denoted by the red "x").

Figure A.6: Robustness Check: Cross-sectional Adaptation to Urban Heat in Alternative Specifications




Notes: Robustness check for figure 1.7. Also includes city-date fixed effects. P-values are calculated under the null hypothesis that the differences in the urban heat effect across high-low groups are all 0.

Figure A.7: Robustness check: Over-Time Adaptation to Urban Heat in Alternative Specifications




Notes: Robustness check for figure 1.7. Also includes city-date fixed effects. P-values are calculated under the null hypothesis that the differences in the urban heat effect across high-low groups are all 0.

Figure A.8: Adaptation to Urban Heat by Previous UH Exposure Levels




Notes: Estimates are derived from Equation 1.4, with UH bins interacting with high-low UH indicators, defined as the average level of urban heat exposure over the past. The p-value tests the null hypothesis that the effects of UH are the same for all high versus low variance groups.

Figure A.9: Placebo Test: Overall Temperature Tomorrow on EMS Caseload Today




Notes: Estimates are based on Equation 1.2, incorporating additional indicators for next-day overall temperature bins with the same set of fixed effects. Coefficients represent next-day overall temperature bins. The p-value tests the null hypothesis that these coefficients are all zero.

Figure A.10: Placebo Test: Urban Heat Tomorrow on EMS Caseload Today




Notes: Estimates are based on Equation 1.4, incorporating additional indicators for next-day urban heat bins with the same set of fixed effects. Coefficients represent next-day urban heat bins. The p-value tests the null hypothesis that these coefficients are all zero.

Figure A.11: Effects of UH on EMS Caseload by % Nonwhite Population




Notes: Estimates are based on Equation 1.4, with city-by-date, calendar-year, and block-group fixed effects. UH bins interact with high-low socioeconomic indicators. Within each panel, results are grouped by season and demographic composition. The p-value tests the null hypothesis that UH effects are identical across high and low share groups.

Figure A.12: Effects of UH on EMS Caseload by % Hispanic Population



Notes: Same setup as Figure B.3, examining the Hispanic population.

Figure A.13: Effects of UH on EMS Caseload by % Population Aged 65+



Notes: Same setup as Figure B.3, examining the population aged 65 and older.

Table A.1: Robustness Check: effects of L in Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)
L <55	-0.0142 (0.0138)	-0.0279 (0.0146)	-0.0273 (0.0144)	-0.0429** (0.0146)	-0.00922 (0.0145)	-0.0103 (0.0147)
[55-60)	-0.0231 (0.0129)	-0.0317* (0.0134)	-0.0311* (0.0133)	-0.0431** (0.0146)	-0.0198 (0.0133)	-0.0150 (0.0131)
[60-65)	0.000697 (0.0107)	-0.00620 (0.0109)	-0.00584 (0.0108)	-0.0241* (0.0112)	0.000158 (0.0112)	0.000635 (0.0114)
[65-70)	-0.00642 (0.00968)	-0.0112 (0.00985)	-0.0109 (0.00980)	-0.0297** (0.00999)	-0.00684 (0.0102)	-0.00842 (0.0104)
[70-75)	-0.00222 (0.00880)	-0.00490 (0.00882)	-0.00469 (0.00881)	-0.0189* (0.00832)	-0.00296 (0.00894)	-0.00245 (0.00889)
[80-85)	0.00385 (0.00780)	0.00831 (0.00795)	0.00861 (0.00792)	0.00862 (0.00889)	0.00876 (0.00811)	0.0122 (0.00815)
[85-90)	0.00158 (0.00859)	0.0104 (0.00891)	0.0110 (0.00884)	0.00837 (0.00955)	0.0160 (0.00919)	0.0204* (0.00940)
[90-95)	-0.00283 (0.00947)	0.0101 (0.00964)	0.0111 (0.00963)	0.0148 (0.0123)	0.0202* (0.00990)	0.0244* (0.0103)
[95-100)	0.00551 (0.0107)	0.0233* (0.0108)	0.0248* (0.0107)	0.0214 (0.0130)	0.0389*** (0.0109)	0.0433*** (0.0111)
[100-105)	0.0209 (0.0124)	0.0416*** (0.0124)	0.0431*** (0.0123)	0.0426** (0.0152)	0.0598*** (0.0128)	0.0646*** (0.0130)
[105-110)	0.0449** (0.0139)	0.0675*** (0.0136)	0.0694*** (0.0135)	0.0748*** (0.0161)	0.0912*** (0.0136)	0.0940*** (0.0137)
[110-115)	0.0425** (0.0151)	0.0674*** (0.0150)	0.0694*** (0.0149)	0.0733*** (0.0175)	0.0971*** (0.0151)	0.0989*** (0.0153)
>115	0.0668*** (0.0175)	0.0932*** (0.0172)	0.0957*** (0.0171)	0.0937*** (0.0206)	0.123*** (0.0167)	0.126*** (0.0166)
UH (5km)		-0.00225*** (0.000602)		-0.00252*** (0.000684)	-0.00179** (0.000565)	-0.00200*** (0.000557)
UH (10km)			-0.00227*** (0.000528)			
UH at 5km-buffer		x		x	x	x
UH at 10km-buffer			x			
N dispatches as Y				x		
BG FE	x	x	x	x		
BG-month FE					x	
BG-week FE						x
N observations	1033081	1033081	1033081	1033081	1032813	1011567

Notes: All specifications also control for calendar year plus day of the year, fixed effects. Standard errors are clustered at the census block group level. Sample sizes are smaller in the last two columns due to additional statistical separation by introducing more flexible fixed effects.

Table A.2: Effects of Overall Temperature L on Health by UH Controls and Season


	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	All	All	May- July	Nov.- Feb.	Rest of the year	All	May- July	Nov.- Feb.	Rest of the year
Overall temperature L	0.00117*** (0.000251)	0.00194*** (0.000298)	0.00208*** (0.000587)	0.00103 (0.000531)	0.00266*** (0.000494)	0.00164*** (0.000288)	0.00184** (0.000567)	0.000720 (0.000531)	0.00220*** (0.000477)
UH		-0.00249*** (0.000523)	-0.00150 (0.00104)	-0.00294** (0.000949)	-0.00328*** (0.000890)				
$UH < -5$						0.0251* (0.0105)	0.0561** (0.0198)	0.00897 (0.0195)	0.0175 (0.0184)
$UH \in [-5, -1)$						0.0123 (0.00793)	0.0455* (0.0180)	-0.00181 (0.0112)	0.0240 (0.0151)
$UH \in [3, 7)$						-0.000685 (0.00524)	0.0158 (0.0143)	-0.00674 (0.00694)	0.00331 (0.0107)
$UH \in [7, 11)$						-0.00994 (0.00627)	0.0261 (0.0146)	-0.0230* (0.00994)	-0.00865 (0.0111)
$UH \in [11, 15)$						-0.0156* (0.00773)	0.0238 (0.0155)	-0.0486* (0.0207)	-0.0183 (0.0127)
$UH > 15$						-0.0102 (0.0102)	0.0340 (0.0184)	-0.000606 (0.0706)	-0.0246 (0.0158)
Continuous L	x	x	x	x	x	x	x	x	x
Continuous UH		x	x	x	x				
Binned UH						x	x	x	x
Joint P-value						0.035	0.048	0.166	0.234
Observations	1033081	1033081	263005	422256	347820	1033081	263005	422256	347820

Notes: Controlling for day-of-the-year, calendar-year, and block-group fixed effects. 7 UH bins in column (3). UH $\in (-1, 3]$ is omitted as the reference group. Standard errors are clustered at the block-group level, reported at 95% in parentheses. P-value tests the null hypothesis that all estimates of urban heat bins are 0.

Appendix B


Appendix for Chapter 2

Figure B.1: Map of Air Quality Changes




Notes: Similar to Figure 2.3, these three panels illustrate changes in average NO_2 and $PM_{2.5}$ concentrations in Northern California and NO_2 in Southern California, before and after the city-level implementation of the minimum wage policy. The analysis spans up to five years before and three years after policy implementation. The data are presented in logarithmic form, de-trended by calendar year, and calculated at the census block group level. The inner (purple) lines delineate the city limits of the minimum wage policy, while the outer (black) lines define the boundaries of the analysis sample. The analysis sample encompasses all census block groups within the same commuting zone and located no more than 50 miles from the policy cities. Areas situated between the two sets of lines are designated as “Outside policy limits” in the analysis.


Figure B.2: Traffic Impact by Time of Day and Day of the Week



Notes: This figure reports the impacts on vehicle miles traveled and average speed by day of the week (y-axis) and hour of the day (x-axis). The regressions are similar to Equation 2.10, but with outcome variables representing the annual average for a given hour on a specific day of the week. I control for location-specific and county-year fixed effects. A larger shade indicates a smaller p-value.


Figure B.3: Heterogeneous Impacts on NO_2 and Vehicle Hours(a) NO_2 

(b) Vehicle Hours Spent





Notes: The top panel reports the policy impacts on ambient NO_2 concentration and vehicle hours by census block group location (within/outside policy limits) $\times$ neighborhood characteristic ranking (top/bottom 50%). Each sub-panel is based on a specific characteristic, following the same order as discussed in the note for Figure 2.4. Characteristic rankings are derived from the 5-year ACS 2010–2014. Outcome variables are in logarithmic form with 95% C.I. reported. Figure notes report the two wild bootstrap p-values that test whether the high versus the low group experiences the impacts.

Figure B.5: Lack of Heterogeneous Impacts on Residential Sorting By Demographic and Socioeconomic Conditions



Notes: This figure reports the policy impacts on resident size by census block group location (within/ outside policy limits) × neighborhood characteristic ranking (top/bottom 50%). Within each sub-panel, estimates are arranged in a similar way as in Figure 2.4, with a detailed description in the note of Figure 2.4. Outcome variables are in logarithm form with 95% C.I. reported. Figure notes report the two wild bootstrap p-values that test whether the high versus the low group experiences the impacts.


Figure B.4: Heterogeneous Impacts on Average Speed and Congestion Probability




Notes: The top panel reports the policy impacts on average vehicle speed and congestion probability by census block group location (within/outside policy limits) $\times$ neighborhood characteristic ranking (top/bottom 50%). Each sub-panel is based on a specific characteristic, following the same order as discussed in the note for Figure 2.4. Characteristic rankings are derived from the 5-year ACS 2010–2014. Outcome variables are in levels with 95% C.I. reported. Figure notes report the two wild bootstrap p-values that test whether the high versus the low group experiences the impacts.

Figure B.6: Policy Impacts by Amenity and Commute Access

(a) By Amenity



(b) By Commute Access



Notes: This figure reports the policy impacts on resident size by census block group location (within/outside policy limits) $\times$ neighborhood amenity or commute access ranking (top/bottom 50%). Within each sub-panel, estimates are arranged in a similar way as in Figure 2.4. 95% confidence intervals are reported with two wild bootstrap p-values that test whether the high versus the low group experiences the impacts in figure notes.

Table B.1: Robustness Check: Policy Impacts with Alternative Fixed Effects

	Within Policy Limits			Outside Policy Limits		
<i>A. Air quality</i>						
$PM_{2.5}$	0.0091 [0.3968]	0.0100 [0.0581]	0.0013 [0.1904]	0.0169 [0.3106]	0.0072 [0.0501]	0.0074* [0.0341]
NO_2	0.0020 [0.4549]	0.0178 [0.2445]	-0.0001 [0.4509]	0.0322 [0.2786]	0.0056* [0.0120]	0.0061* [0.0100]
<i>B. Traffic</i>						
Vehicle Mile Traveled	-0.0053 [0.3186]	0.0357 [0.2164]	0.0034 [0.3206]	0.0265* [0.0140]	0.0374* [0.0100]	0.0330** [0.0080]
Vehicle Hour Spent	-0.0167 [0.0842]	0.0523 [0.2164]	0.0077 [0.2485]	0.0083 [0.1643]	0.0421* [0.0120]	0.0384* [0.0140]
Average Speed	-0.0115 [0.4609]	-0.5295 [0.0541]	-0.1819* [0.0240]	-0.1503* [0.0481]	-0.1323* [0.0160]	-0.1567* [0.0180]
% Congestion	0.0024 [0.3647]	0.0339** [0.0040]	0.0085 [0.1122]	0.0018 [0.2886]	0.0070* [0.0140]	0.0073** [0.0020]
% Severe Congestion	0.0004 [0.4309]	0.0128 [0.0521]	0.0042* [0.0120]	0.0025 [0.1443]	0.0042* [0.0100]	0.0039* [0.0120]
<i>C. Commute</i>						
Emp-Res Ratio	0.0162 [0.1764]	0.0102 [0.3928]	0.0101 [0.1122]	-0.0165* [0.0441]	-0.0276 [0.0661]	-0.0153 [0.2104]
Median Driving Distance	0.0729 [0.2425]	0.0472 [0.3347]	0.0589 [0.2585]	0.0916 [0.2886]	-0.3412 [0.2786]	-0.2782 [0.1924]
Median Driving Minutes	0.0490 [0.2725]	0.0618 [0.2325]	0.0305 [0.2605]	0.0795 [0.2325]	-0.2361 [0.2786]	-0.2086 [0.2806]
County-Year FE			x			x
County-Linear	x			x		
County-Sub-Div. FE		x			x	
N	42698	42698	42698	80189	80189	80189

Notes: This table reports policy effects using alternative specifications. The benchmark model (columns 3 and 6) includes county-by-year fixed effects. The other columns use county-specific linear trends or county-subdivision-by-year fixed effects. Estimations are based on equation 2.10. Standard errors are clustered at the policy city level to account for arbitrary correlation within the cluster.

Table B.2: Policy Impacts Accounting for Leads and Lags

	Vehicle Miles	Vehicle Hours	Avg. Speed	Congestion	$PM_{2.5}$	NO_2	Emp-Res Ratio	Driving Minutes	Driving Miles
<i>Panel A. Inside Policy Limits, N = 42698</i>									
MW	0.0196	0.0189	0.0383	0.000260	0.000220	0.0103	0.0122	0.0714	0.0638
MW, 1-year lag	0.00264	0.000769	0.0967	0.000740	-0.000943	0.00238	-0.0115	0.0201	0.0190
MW, 2-year lag	0.0251	0.0259	0.0221	0.00135	0.00561	0.00798	-0.0127	0.153	0.122
MW, 2-year lead	0.0178	0.0175	0.190	-0.00999	-0.00127	0.00318	-0.0260	0.0758	0.0624
MW, 1-year lead	-0.0226	-0.0174	-0.429	0.0136	0.00115	-0.0127	0.0132	0.0184	0.0272
P-value of MW	0.166	0.200	0.246	0.457	0.439	0.106	0.178	0.016	0.008
Joint P-value of Leads	0.059	0.264	0.024	0.110	0.424	0.188	0.016	0.404	0.417
Joint P-value of Lags	0.048	0.061	0.445	0.497	0.084	0.359	0.376	0.090	0.085
<i>Panel B. Outside Policy Limits, N = 80189</i>									
MW	0.0215	0.0254	-0.148	0.00595	0.00214	0.0119	-0.00941	-0.196	-0.130
MW, 1-year lag	-0.00314	-0.00488	0.0779	0.000208	0.00695	-0.00464	-0.00364	-0.206	-0.187
MW, 2-year lag	0.0523	0.0485	-0.0392	0.00542	-0.0175	0.0157	-0.00481	-0.220	-0.128
MW, 2-year lead	-0.0336	-0.0373	0.155	-0.00294	-0.00845	-0.00125	-0.0456	-0.256	-0.213
MW, 1-year lead	0.0569	0.0617	-0.147	0.00758	0.00140	-0.00739	-0.00222	-0.354	-0.260
P-value of CMW	0.022	0.036	0.076	0.062	0.016	0.088	0.212	0.008	0.104
Joint P-value of Leads	0.082	0.061	0.079	0.163	0.139	0.078	0.236	0.136	0.077
Joint P-value of Lags	0.284	0.287	0.156	0.156	0.015	0.175	0.499	0.011	0.024

Notes: This table is based on equation 2.10, additionally accounting for two-year leads and lags. I do not indicate significance levels in the table. Instead, I report wild clustered bootstrap p-values for the contemporaneous MW term, and joint p-values for leads and lags.

Table B.3: Robustness Check: Policy Impacts using Heterogeneity-Robust Estimator

	Within the Policy Limits		Outside the Policy Limits	
	Main Effect	Placebo Effect	Main Effect	Placebo Effect
<i>Panel A. Air quality</i>				
$PM_{2.5}$	0.0138 (0.0131)	0.0171 (0.0378)	0.0371*** (0.0062)	-0.0123 (0.0308)
NO_2	0.0657*** (0.0144)	0.1013 (0.1027)	0.0976*** (0.0142)	0.1143 (0.0829)
<i>Panel B. Traffic</i>				
Vehicle Miles	-0.0198 (0.0141)	0.0487* (0.0234)	0.0624*** (0.0138)	-0.0192 (0.0146)
Vehicle Hours	-0.0074 (0.0131)	0.0241 (0.0261)	0.0620*** (0.0114)	-0.0300 (0.0218)
Average Speed	-0.0530 (0.0826)	0.8534** (0.3225)	-0.0682 (0.0462)	0.1620 (0.1180)
<i>Panel C. Commute</i>				
Emp-Res Ratio	0.0109*** (0.0029)	-0.0043 (0.0053)	-0.0023 (0.0092)	0.0098* (0.0044)
Driving Miles	0.0941 (0.1256)	0.0661 (0.0463)	-0.0732** (0.0248)	-0.0935 (0.2708)
Driving Minutes	0.3221** (0.1172)	0.1381 (0.0808)	-0.0641* (0.0261)	-0.0038 (0.2083)
N	23391	23391	43652	43652

Notes: This table reports heterogeneous treatment effects estimated using the method of de Chaisemartin et al. (2022), which allows for non-linearity. I use normalized minimum wage levels as the policy variable to mirror the main specification and compare cities with relatively high versus low policy intensity. Normalized minimum wage refers to the difference between city- and state-level minimum wage. I include policy leads for vehicle miles and time outcomes to account for anticipation effects, consistent with Appendix Table B.2. The placebo test compares switchers and stayers during the period before the switch occurs. Standard errors are clustered at the policy city level to account for within-cluster correlation.

Table B.4: Robustness Check: Policy Impacts using Average Minimum Wage Levels as Policy Variable

	Within Policy Limits	Outside Policy Limits
<i>Panel A. Air quality</i>		
PM 2.5	0.0115 [0.0804]	0.0025 [0.4422]
NO2	0.0013 [0.4020]	0.0092* [0.0302]
<i>Panel B. Traffic</i>		
Vehicle Miles	-0.0398 [0.4322]	0.0118 [0.7940]
Vehicle Hours	-0.0327 [0.4925]	0.0213 [0.7739]
Average Speed	-0.3039 [0.0603]	-0.3083 [0.0503]
Congestion	0.0103 [0.3216]	0.0126* [0.0201]
<i>Panel C. Commute</i>		
Emp-Res Ratio	0.0301* [0.0402]	-0.0154 [0.1206]
Driving Miles	0.3045*** [0.0000]	-0.0923 [0.5528]
Driving Minutes	0.3858*** [0.0000]	-0.1591 [0.4221]
N	42698	80189

Notes: This table reports the policy effects using an alternative policy variable, defined as the average of minimum wage levels across nearby policy cities. Estimations are based on equation 2.10, with **wild clustered bootstrap P-values** reported in square brackets.

Table B.5: Robustness Check: Policy Impacts in Stacked Design using Continuous Treatment

	Within Policy Limits	Outside Policy Limits
<i>Panel A. Air quality</i>		
$PM_{2.5}$	0.01086 (0.00624)	0.01379 (0.01113)
NO_2	0.06474*** (0.00770)	0.09906* (0.04030)
<i>Panel B. Traffic</i>		
Vehicle Miles	0.00847 (0.01264)	0.01961 (0.02284)
Vehicle Hours	0.01345 (0.01351)	0.02162 (0.02196)
Average Speed	-0.19594*** (0.03935)	-0.03558 (0.05873)
<i>Panel C. Commute</i>		
Emp-Res Ratio	0.01490 (0.00954)	-0.01424 (0.01132)
Driving Miles	-0.00272 (0.04086)	0.01332 (0.05712)
Driving Minutes	0.01079 (0.04918)	0.05121 (0.06826)
N	330398	483700

~~Notes: This table reports stacked estimates following Congiz et al. (2019) and Dube & Lindner (2024).~~
The comparison is between a city with a minimum wage and cities that adopt the policy in later years (but not concurrently). Each comparison is treated as an event, and events are stacked into a new data frame for estimation. We control for event-location and event-time fixed effects so that the variation comes solely from within each event. Standard errors are clustered at the policy city level to account for arbitrary correlation within the cluster.

Table B.6: Robustness Check: Policy Impacts in Stacked Design using Binary Treatment

	Within Policy Limits	Outside Policy Limits
<i>Panel A. Air quality</i>		
$PM_{2.5}$	0.00209 (0.02650)	0.00121 (0.02417)
NO_2	0.08953 (0.05829)	0.03159 (0.10334)
<i>Panel B. Traffic</i>		
Vehicle Miles	0.01473 (0.02661)	0.01322 (0.05948)
Vehicle Hours	0.02509 (0.02792)	0.01809 (0.05855)
Average Speed	-0.40932** (0.13018)	-0.11367 (0.07307)
<i>Panel C. Commute</i>		
Emp-Res Ratio	0.07269 (0.04144)	-0.01136 (0.03596)
Driving Miles	-0.02871 (0.07027)	0.08075 (0.07954)
Driving Minutes	0.00238 (0.08505)	0.17337 (0.09073)
N	330398	483700

Notes: This table reports stacked estimates following Cengiz et al. (2019) and Dube & Lindner (2024). The comparison is between a city with a minimum wage and cities that adopt the policy in later years (but not concurrently). Each comparison is treated as an event, and events are stacked into a new data frame for estimation. We control for event-location and event-time fixed effects so that the variation comes solely from within each event. Standard errors are clustered at the policy city level to account for arbitrary correlation within the cluster.

Table B.7: Policy Impacts on Public Transportation Ridership

	(1)	(2)	(3)	(4)
	All	Bus	Rail	Others
Minimum Wage	-0.00186	-0.0352	0.00246	-0.00575
	(0.0215)	(0.0257)	(0.00880)	(0.0159)

N 314 306 314 314
Notes: The outcome variable is log ridership from the Federal Transit Administration's (FTA) National Transit Database (NTD) at the transit agency level each year. Due to the coarse geographical level of the data, these analyses do not differentiate agencies within versus outside policy limits. Each transit agency is assigned to the nearest policy city. We control for transit-agency-specific and urbanized-area-by-year fixed effects, with standard errors clustered at the transit-agency level.

Table B.8: Air Quality Impacts for Near-Road and Far-Road Block Groups

	All Sample	Top 10%	Top 25%	Bottom 25%
<i>Panel A. Inside Policy Limits</i>				
$PM_{2.5}$	0.0013	0.0023	0.0002	0.0043
	[0.1904]	[0.2164]	[0.4349]	[0.0962]
NO_2	-0.0001	-0.0038	-0.0003	0.0003
	[0.4509]	[0.3226]	[0.4729]	[0.3707]
<i>N</i>	42,698	4,273	10,722	10,627
<i>Panel B. Outside Policy Limits</i>				
$PM_{2.5}$	0.0074*	0.0108	0.0074	0.0049
	[0.0341]	[0.0641]	[0.0601]	[0.0822]
NO_2	0.0061*	0.0119*	0.0065*	-0.0026
	[0.0100]	[0.0160]	[0.0200]	[0.2986]

N 80,189 7,905 19,785 21,159
Notes: This table repeats the air quality impacts estimation for different samples. The first column represents the whole sample as shown in Table 2.3. The second to the fourth columns use sub-samples of block groups defined by their ranking in total length of roads within a 1-mile radius. That is, the second row reports the impacts for block groups that rank in the top 10% in terms of the total nearby highway lengths, while the third and fourth columns report results for the top and bottom 25%. Wild clustered bootstrap p-values are reported in square brackets.

Appendix C

Appendix for Chapter 3

Table C.1: Variable means by generation and sex

	NLSY 79 women (1)	NLSY 79 Men (2)	NLSY 97 Women (3)	NLSY 97 Men (4)
Panel A: Age-invariant variables				
High-risk	0.484	0.487	0.408	0.428
Mom's highest grade completed	10.584	10.678	11.807	11.850
Dad's highest grade completed	10.122	10.262	11.612	11.550
Family size	5.490	5.555	4.683	4.672
Black	0.471	0.479	0.435	0.425
Lived with both parents at 12/14	0.574	0.590	0.258	0.284
Return to education, age 18-22	0.119	0.119	0.133	0.133
Incarceration rate, age 18-22	187.315	186.819	455.827	451.812
The income shares for the top 10%, age 18-22	36.492	36.502	43.977	44.165
Real GDP, age 18-22	29.911	30.160	47.814	47.722
Unemployment rate, age 18-22	8.467	8.510	5.213	5.233
Highest grade completed	12.342	11.987	13.022	12.315
High school graduation	0.863	0.792	0.870	0.836
Sub-baccalaureate credential	0.181	0.142	0.243	0.155
BA or greater	0.099	0.091	0.172	0.105
<i>N</i>	1,378	1,448	2,095	2,137
Panel B: Age-variant variables				
Any employment	0.760	0.818	0.851	0.829
Weeks worked	27.946	30.772	34.754	34.203
Hours worked	989	1,223	1,263	1,428
Working, enrolled	0.125	0.115	0.309	0.233
Working, not enrolled	0.402	0.451	0.542	0.592
Not working, enrolled	0.116	0.126	0.035	0.040
Not working, not enrolled	0.356	0.307	0.113	0.134
Working, with children	0.098	0.074	0.228	0.114
Working, without children	0.428	0.491	0.627	0.712
Enrolled, with children	0.022	0.009	0.040	0.010
Enrolled, without children	0.219	0.232	0.314	0.269
<i>N</i>	6,891	7,242	10,477	10,683

Table C.2: Risk regressions, by estimation sample and function form

A. Ordinary least squares				
	(1) $Avg. \frac{income}{needs}$	(2) $\frac{Avg.income}{Avg.needs}$	(3) $Avg. \frac{income}{needs}$	(4) $\frac{Avg.income}{Avg.needs}$
Mom's HGC	-0.085 (0.064)	-0.084 (0.064)	-0.096 (0.059)	-0.095 (0.059)
Dad's HGC	0.046 (0.055)	0.049 (0.055)	0.038 (0.051)	0.040 (0.051)
Family size	-0.163 (0.036)	-0.162 (0.035)	-0.115 (0.039)	-0.114 (0.039)
Mom's HGC ²	0.011 (0.003)	0.011 (0.003)	0.012 (0.003)	0.012 (0.003)
Dad's HGC ²	0.008 (0.002)	0.008 (0.002)	0.008 (0.002)	0.007 (0.002)
Mom's HGC $\times$ Dad's HGC	-0.003 (0.002)	-0.003 (0.002)	-0.005 (0.002)	-0.005 (0.002)
Mom's HGC $\times$ Family size	0.005 (0.004)	0.005 (0.004)	0.003 (0.004)	0.003 (0.004)
Dad's HGC $\times$ Family size	-0.008 (0.003)	-0.008 (0.003)	-0.008 (0.003)	-0.008 (0.003)
Lived with both parents at 12/14			1.024 (0.256)	1.031 (0.255)
Lived with both parents at 12/14 $\times$ Mom's HGC			0.019 (0.020)	0.019 (0.020)
Lived with both parents at 12/14 $\times$ Dad's HGC			0.036 (0.016)	0.036 (0.016)
Lived with both parents at 12/14 $\times$ Family size			-0.119 (0.025)	-0.119 (0.025)
Constant	1.707 (0.385)	1.686 (0.382)	1.261 (0.378)	1.236 (0.373)
Observations	3182	3182	3182	3182
R^2	0.330	0.330	0.400	0.401

Notes: Standard errors are in parentheses, clustered by household.

B. Weighted least squares				
	(1) <i>Avg. $\frac{income}{needs}$</i>	(2) <i>Avg. $\frac{income}{needs}$</i>	(3) <i>Avg. $\frac{income}{needs}$</i>	(4) <i>Avg. $\frac{income}{needs}$</i>
Mom's HGC	-0.089 (0.082)	-0.087 (0.082)	-0.107 (0.075)	-0.104 (0.075)
Dad's HGC	0.145 (0.079)	0.149 (0.078)	0.119 (0.073)	0.121 (0.073)
Family size	-0.243 (0.054)	-0.243 (0.053)	-0.183 (0.057)	-0.182 (0.056)
Mom's HGC ²	0.011 (0.004)	0.011 (0.004)	0.012 (0.004)	0.012 (0.004)
Dad's HGC ²	0.006 (0.003)	0.006 (0.003)	0.006 (0.003)	0.006 (0.003)
Mom's HGC $\times$ Dad's HGC	-0.007 (0.003)	-0.007 (0.003)	-0.008 (0.002)	-0.008 (0.002)
Mom's HGC $\times$ Family size	0.011 (0.006)	0.011 (0.006)	0.008 (0.006)	0.008 (0.006)
Dad's HGC $\times$ Family size	-0.008 (0.004)	-0.008 (0.004)	-0.007 (0.004)	-0.007 (0.004)
Lived with both parents at 12/14			0.830 (0.349)	0.856 (0.346)
Lived with both parents at 12/14 $\times$ Mom's HGC			0.040 (0.027)	0.039 (0.026)
Lived with both parents at 12/14 $\times$ Dad's HGC			0.022 (0.022)	0.022 (0.022)
Lived with both parents at 12/14 $\times$ Fam- ily size			-0.118 (0.035)	-0.120 (0.035)
Constant	1.690 (0.537)	1.657 (0.532)	1.341 (0.502)	1.298 (0.494)
Observations	3182	3182	3182	3182

Notes: Standard errors are clustered by household. Weights are wave-1 sampling weights.

Appendix D

Permissions and Attributions

Chapter 2 was previously published in the Journal of Environmental Economics and Management (Bian 2025). It is reproduced here in accordance with the publisher Elsevier's licensing and copyright policies: <https://www.elsevier.com/en-au/about/policies-and-standards/copyright#1-author-rights>.

Chapter 3 was co-authored with Jeffrey Grogger at The University of Chicago and Jacob Bastian at Rutgers University, and was previously published in the Labour Economics (Bastian et al. 2022). It is reproduced here in accordance with the publisher Elsevier's licensing and copyright policies: <https://www.elsevier.com/en-au/about/policies-and-standards/copyright#1-author-rights>.

Bibliography

- Albouy, D. (2016), ‘What are cities worth? land rents, local productivity, and the total value of amenities’, *Review of Economics and Statistics* **98**(3), 477–487.
- Allegretto, S. A., Dube, A. & Reich, M. (2011), ‘Do minimum wages really reduce teen employment? accounting for heterogeneity and selectivity in state panel data’, *Industrial Relations: A Journal of Economy and Society* **50**(2), 205–240.
- Allegretto, S. & Reich, M. (2018), ‘Are local minimum wages absorbed by price increases? estimates from internet-based restaurant menus’, *ILR Review* **71**(1), 35–63.
- Anderson, M. (2017), ‘How to beat the traffic’, *Science* **357**(6346), 36–37.
- Anderson, M. L. (2020), ‘As the wind blows: The effects of long-term exposure to air pollution on mortality’, *Journal of the European Economic Association* **18**(4), 1886–1927.
- Barreca, A., Clay, K., Deschênes, O., Greenstone, M. & Shapiro, J. S. (2015), ‘Convergence in adaptation to climate change: Evidence from high temperatures and mortality, 1900–2004’, *American Economic Review* **105**(5), 247–251.
- Barreca, A., Clay, K., Deschênes, O., Greenstone, M. & Shapiro, J. S. (2016), ‘Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century’, *Journal of Political Economy* **124**(1), 105–159.
- Barreca, A. I. (2012), ‘Climate change, humidity, and mortality in the united states’, *Journal of Environmental Economics and Management* **63**(1), 19–34.
- Bastian, J. (2020), ‘The rise of working mothers and the 1975 earned income tax credit’, *American Economic Journal: Economic Policy* **12**(3), 44–75.
- Bastian, J., Bian, L. & Grogger, J. (2022), ‘How did safety-net reform affect the education of adolescents from low-income families?’, *Labour Economics* **77**, 102031.
- Bastian, J., Bian, L. & Grogger, J. (forthcoming), ‘How did safety-net reform affect early adulthood among adolescents from low-income families?’, *National Tax Journal* .

BIBLIOGRAPHY

- Bastian, J. E. & Jones, M. R. (2020), Do EITC Expansions Pay for Themselves? Effects on Tax Revenue and Public Assistance Spending, Technical report, Working Paper.
URL: <https://nta.confex.com/nta/2018/meetingapp.cgi/Paper/1427>
- Bastian, J. & Lochner, L. (2020), The eitc and maternal time use: More time working and less time with kids?, Technical report, National Bureau of Economic Research.
- Bastian, J. & Micheltore, K. (2018), ‘The Long-Term Impact of the Earned Income Tax Credit on Children’s Education and Employment Outcomes’, *Journal of Labor Economics* **36**(4), 1127–1163.
- Bavier, R. (2001), ‘Welfare reform data from the survey of income and program participation’, *Monthly Lab. Rev.* **124**, 13.
- Becker, G. S. & Murphy, K. M. (1988), ‘A theory of rational addiction’, *Journal of political Economy* **96**(4), 675–700.
- Benali, A., Carvalho, A., Nunes, J., Carvalhais, N. & Santos, A. (2012), ‘Estimating air surface temperature in portugal using modis lst data’, *Remote Sensing of Environment* **124**, 108–121.
- Benz, S. A. & Burney, J. A. (2021), ‘Widespread race and class disparities in surface urban heat extremes across the united states’, *Earth’s Future* **9**(7), e2021EF002016.
- Bian, L. (2025), ‘Do microclimates matter? the health impacts of urban heat using fine-scale data’, *Journal of the Association of Environmental and Resource Economists* **13**, 2025–ics0095069625000750.
URL: <https://ideas.repec.org/a/eee/jeeman/v133y2025ics0095069625000750.html>
- Blank, R. M. (2002), ‘Evaluating welfare reform in the united states’, *Journal of Economic Literature* **40**(4), 1105–1166.
- Borg, G., Passaro, D. G. & Herno, S. (2022), ‘From workplace to residence: The spillover effects of minimum wage policies on local housing markets’, *arXiv preprint arXiv:2208.01791* .
- Borsky, S., Fesselmeier, E. & Vogelsang, L. (2024), ‘Urban heat and within-city residential sorting’, *Journal of Environmental Economics and Management* p. 103014.
- Boudreaux, M. H., Golberstein, E. & McAlpine, D. D. (2016), ‘The long-term impacts of medicaid exposure in early childhood: Evidence from the program’s origin’, *Journal of health economics* **45**, 161–175.
- Brack, C. L. (2002), ‘Pollution mitigation and carbon sequestration by an urban forest’, *Environmental pollution* **116**, S195–S200.

BIBLIOGRAPHY

- Brown, D. W., Kowalski, A. E. & Lurie, I. Z. (2020), ‘Long-Term Impacts of Childhood Medicaid Expansions on Outcomes in Adulthood’, *Review of Economic Studies* **87**(2), 792–821.
- Buchmueller, T., Ham, J. C. & Shore-Sheppard, L. D. (2016), The Earned Income Tax Credit, in ‘Economics of Means-Tested Transfer Programs in the United States, Volume I’, Vol. I, University of Chicago Press, pp. 1–119.
- Burnett, R., Chen, H., Szyszkowicz, M., Fann, N., Hubbell, B., Pope III, C. A., Apte, J. S., Brauer, M., Cohen, A., Weichenthal, S. et al. (2018), ‘Global estimates of mortality associated with long-term exposure to outdoor fine particulate matter’, *Proceedings of the National Academy of Sciences* **115**(38), 9592–9597.
- Busso, M., Gregory, J. & Kline, P. (2013), ‘Assessing the incidence and efficiency of a prominent place based policy’, *American Economic Review* **103**(2), 897–947.
- Cadena, B. C. (2014), ‘Recent immigrants as labor market arbitrageurs: Evidence from the minimum wage’, *Journal of urban Economics* **80**, 1–12.
- California State Legislature (2024), ‘California 2019 Low-Wage Workers and Minimum Wage — lao.ca.gov’, <https://lao.ca.gov/Publications/Report/4878/1>. [Accessed 25-08-2025].
- Card, D. & Krueger, A. B. (1994), ‘Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania’, *The American Economic Review* **84**(4), 772.
- Carleton, T. A. (2017), ‘Crop-damaging temperatures increase suicide rates in india’, *Proceedings of the National Academy of Sciences* **114**(33), 8746–8751.
- Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Kopp, R. E., McCusker, K. E., Nath, I. et al. (2022), ‘Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits’, *The Quarterly Journal of Economics* **137**(4), 2037–2105.
- Carozzi, F. & Roth, S. (2023), ‘Dirty density: Air quality and the density of american cities’, *Journal of Environmental Economics and Management* **118**, 102767.
- Cengiz, D., Dube, A., Lindner, A. & Zipperer, B. (2019), ‘The effect of minimum wages on low-wage jobs’, *The Quarterly Journal of Economics* **134**(3), 1405–1454.
- Chakma, T., Colmer, J. & Voorheis, J. (2024), The causes and consequences of urban heat islands, in ‘ASHEcon 2024 Conference’, ASHECON.
- Chauhan, A., Inskip, H. M., Linaker, C. H., Smith, S., Schreiber, J., Johnston, S. L. & Holgate, S. T. (2003), ‘Personal exposure to nitrogen dioxide (no₂) and the severity of virus-induced asthma in children’, *The Lancet* **361**(9373), 1939–1944.

BIBLIOGRAPHY

- Chen, Y. & Whalley, A. (2012), ‘Green infrastructure: The effects of urban rail transit on air quality’, *American Economic Journal: Economic Policy* **4**(1), 58–97.
- Chetty, R., Friedman, J. N. & Rockoff, J. (2011), ‘New Evidence on the Long-Term Impacts of Tax Credits’, *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association* **104**, 1–43.
- Clinton, N. & Gong, P. (2013), ‘Modis detected surface urban heat islands and sinks: Global locations and controls’, *Remote Sensing of Environment* **134**, 294–304.
- Cohodes, S. R., Grossman, D. S., Kleiner, S. A. & Lovenheim, M. F. (2016), ‘The Effect of Child Health Insurance Access on Schooling: Evidence from Public Insurance Expansions’, *Journal of Human Resources* **51**(3), 727–759.
- Cook, C. (2025), ‘Heterogeneous preferences for neighborhood amenities: Evidence from gps data’, *Review of Economics and Statistics* pp. 1–43.
- Cooper, M. J., Martin, R. V., McLinden, C. A. & Brook, J. R. (2020), ‘Inferring ground-level nitrogen dioxide concentrations at fine spatial resolution applied to the tropomi satellite instrument’, *Environmental Research Letters* **15**(10), 104013.
- Couture, V., Gaubert, C., Handbury, J. & Hurst, E. (2024), ‘Income growth and the distributional effects of urban spatial sorting’, *Review of Economic Studies* **91**(2), 858–898.
- Crouse, D. L., Peters, P. A., Hystad, P., Brook, J. R., van Donkelaar, A., Martin, R. V., Villeneuve, P. J., Jerrett, M., Goldberg, M. S., Pope III, C. A. et al. (2015), ‘Ambient pm_{2.5}, o₃, and no₂ exposures and associations with mortality over 16 years of follow-up in the canadian census health and environment cohort (canchec)’, *Environmental health perspectives* **123**(11), 1180–1186.
- Currie, J., Voorheis, J. & Walker, R. (2023), ‘What caused racial disparities in particulate exposure to fall? new evidence from the clean air act and satellite-based measures of air quality’, *American Economic Review* **113**(1), 71–97.
- Currie, J. & Walker, R. (2011), ‘Traffic congestion and infant health: Evidence from e-zpass’, *American Economic Journal: Applied Economics* **3**(1), 65–90.
- Dave, D., Corman, H., Kalil, A., Schwartz-Soicher, O. & Reichman, N. E. (2020), ‘Intergenerational Effects of Welfare Reform: Adolescent Delinquent and Risky Behaviors’, *Economic Inquiry* .
- Dave, D. M., Corman, H. & Reichman, N. E. (2012), ‘Effects of Welfare Reform on Education Acquisition of Adult Women’, *Journal of Labor Research* **33**(2), 251–282.

BIBLIOGRAPHY

- de Chaisemartin, C., d'Haultfoeuille, X., Pasquier, F. & Vazquez-Bare, G. (2022), 'Difference-in-differences for continuous treatments and instruments with stayers', *Available at SSRN 4011782*.
- Deschênes, O. (2014), 'Temperature, human health, and adaptation: A review of the empirical literature', *Energy Economics* **46**, 606–619.
- Deschênes, O. & Greenstone, M. (2011), 'Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us', *American Economic Journal: Applied Economics* **3**(4), 152–185.
- Deschênes, O., Greenstone, M. & Shapiro, J. S. (2017), 'Defensive investments and the demand for air quality: Evidence from the nox budget program', *American Economic Review* **107**(10), 2958–2989.
- Deschênes, O. & Moretti, E. (2009), 'Extreme weather events, mortality, and migration', *The Review of Economics and Statistics* **91**(4), 659–681.
- Dharmasankar, S. & Yoo, H. (2023), 'Assessing the main and spillover effects of seattle's minimum wage on establishment decisions', *Regional Science and Urban Economics* **99**, 103865.
- Dube, A. (2019), 'Minimum wages and the distribution of family incomes', *American Economic Journal: Applied Economics* **11**(4), 268–304.
- Dube, A., Girardi, D., Jorda, O. & Taylor, A. M. (2023), A local projections approach to difference-in-differences event studies, Technical report, National Bureau of Economic Research.
- Dube, A., Lester, T. W. & Reich, M. (2010), 'Minimum wage effects across state borders: Estimates using contiguous counties', *The review of economics and statistics* **92**(4), 945–964.
- Dube, A., Lester, T. W. & Reich, M. (2016), 'Minimum wage shocks, employment flows, and labor market frictions', *Journal of Labor Economics* **34**(3), 663–704.
- Dube, A. & Lindner, A. (2021), 'City limits: What do local-area minimum wages do?', *Journal of Economic Perspectives* **35**(1), 27–50.
- Dube, A. & Lindner, A. (2024), 'Minimum wages in the 21st century', *Handbook of Labor Economics* **5**, 261–383.
- Dube, A., Naidu, S. & Reich, M. (2007), 'The economic effects of a citywide minimum wage', *ILR Review* **60**(4), 522–543.
- Duncan, B., Grogger, J., Leon, A. S. & Trejo, S. J. (2020), 'New evidence of generational progress for Mexican Americans', *Labour Economics* **62**, 101771.

BIBLIOGRAPHY

- Duranton, G. & Venables, A. J. (2020), Place-based policies: principles and developing country applications, *in* ‘Handbook of regional science’, Springer, pp. 1–22.
- Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M. & Vom Berge, P. (2022), ‘Reallocation effects of the minimum wage’, *The Quarterly Journal of Economics* **137**(1), 267–328.
- Eissa, N. & Hoynes, H. W. (2004), ‘Taxes and the labor market participation of married couples: The earned income tax credit’, *Journal of Public Economics* **88**(9-10), 1931–1958.
- Faggio, G. (2019), ‘Relocation of public sector workers: Evaluating a place-based policy’, *Journal of Urban Economics* **111**, 53–75.
- Fan, J., Li, S., Fan, C., Bai, Z. & Yang, K. (2016), ‘The impact of pm2. 5 on asthma emergency department visits: a systematic review and meta-analysis’, *Environmental Science and Pollution Research* **23**, 843–850.
- Fan, M., Tang, Q., Wu, J. & Xiao, J. (2023), ‘Unintended environmental consequences of place-based economic policies’, *American Journal of Agricultural Economics* **105**(1), 291–315.
- Feng, S., Gao, D., Liao, F., Zhou, F. & Wang, X. (2016), ‘The health effects of ambient pm2. 5 and potential mechanisms’, *Ecotoxicology and environmental safety* **128**, 67–74.
- Feng, Y., Wang, L. & Nie, C. (2025), ‘Can place-based policy reduce carbon emissions? evidence from industrial transformation and upgrading exemplary zone in china’, *Humanities and Social Sciences Communications* **11**(1), 877.
URL: <https://www.nature.com/articles/s41599-024-03383-w>
- Fretz, S., Parchet, R. & Robert-Nicoud, F. (2022), ‘Highways, market access and spatial sorting’, *The Economic Journal* **132**(643), 1011–1036.
- Gaigné, C., Koster, H. R., Moizeau, F. & Thisse, J.-F. (2022), ‘Who lives where in the city? amenities, commuting and income sorting’, *Journal of Urban Economics* **128**, 103394.
- Gaigné, C., Riou, S. & Thisse, J.-F. (2012), ‘Are compact cities environmentally friendly?’, *Journal of Urban Economics* **72**(2-3), 123–136.
- Ganong, P. & Liebman, J. B. (2018), ‘The Decline, Rebound, and Further Rise in SNAP enrollment: Disentangling Business Cycle Fluctuations and Policy Changes’, *American Economic Journal: Economic Policy* **10**(4), 153–176.
- Gaubert, C. & Robert-Nicoud, F. (2025), Sorting to expensive cities, Technical report, National Bureau of Economic Research.

BIBLIOGRAPHY

- Gennetian, L. A., Duncan, G. J., Knox, V. W., Vargas, W. G., Clark-Kauffman, E. & London, A. S. (2002), How welfare and work policies for parents affect adolescents: a synthesis of research, Technical report, MDRC.
- Glaeser, E. L. (2008), *Cities, agglomeration, and spatial equilibrium*, OUP Oxford.
- Glaeser, E. L. & Kahn, M. E. (2010), ‘The greenness of cities: Carbon dioxide emissions and urban development’, *Journal of urban economics* **67**(3), 404–418.
- Goldin, J. & Micheltore, K. (2020), Who benefits from the child tax credit?, Technical report, National Bureau of Economic Research.
- Goodman-Bacon, A. (2016), The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes, Technical report, National Bureau of Economic Research.
- Gould, C. F., Heft-Neal, S., Heaney, A. K., Bendavid, E., Callahan, C. W., Kiang, M., Zivin, J. S. G. & Burke, M. (2024), Temperature extremes impact mortality and morbidity differently, Technical report, National Bureau of Economic Research.
- Grogger, J. (2003), ‘The effects of time limits, the EITC, and other policy changes on welfare use, work, and income among female-headed families’, *Review of Economics and Statistics* **85**(2).
- Grogger, J. & Karoly, L. (2005), ‘Employment, labor supply, and earnings’, *Welfare reform: Effects of a decade of change* pp. 134–154.
- Grossman, D. (2019), ‘The unintended effects of place based programs: fertility and health effects of urban empowerment zones’, *Journal of health economics* **63**, 114–127.
- Grossman, M. (1972), ‘On the concept of health capital and the demand for health’, *Journal of Political economy* **80**(2), 223–255.
- Gruber, J. (2003), Medicaid, in R. A. Moffitt, ed., ‘Means-tested transfer programs in the United States’, University of Chicago Press, pp. 15–78.
- Hamra, G. B., Laden, F., Cohen, A. J., Raaschou-Nielsen, O., Brauer, M. & Loomis, D. (2015), ‘Lung cancer and exposure to nitrogen dioxide and traffic: a systematic review and meta-analysis’, *Environmental health perspectives* **123**(11), 1107–1112.
- Han, D., Yu, D. & Qiu, J. (2023), ‘Assessing coupling interactions in a safe and just operating space for regional sustainability’, *Nature Communications* **14**(1), 1369.
- Hansen, J., Ruedy, R., Sato, M. & Lo, K. (2010), ‘Global surface temperature change’, *Reviews of Geophysics* **48**(4), 14288–14293.

BIBLIOGRAPHY

- Haskins, R. (2006), *Work over Welfare: The Inside Story of the 1996 Welfare Reform Law*, Brookings Institution Press, Washington, DC.
- Heaviside, C., Macintyre, H. & Vardoulakis, S. (2017), ‘The urban heat island: implications for health in a changing environment’, *Current environmental health reports* **4**, 296–305.
- Heaviside, C., Vardoulakis, S. & Cai, X.-M. (2016), ‘Attribution of mortality to the urban heat island during heatwaves in the west midlands, uk’, *Environmental health* **15**, 49–59.
- Heckman, J., Lochner, L. & Cossa, R. (2002), Learning-by-doing vs. on-the-job training: Using variation induced by the eitc to distinguish between models of skill formation, Technical report, National Bureau of Economic Research.
- Heft-Neal, S., Lobell, D. B. & Burke, M. (2017), ‘Using remotely sensed temperature to estimate climate response functions’, *Environmental Research Letters* **12**(1), 014013.
- Hernandez-Cortes, D., Meng, K. C. & Weber, P. (2023), ‘Decomposing trends in us air pollution disparities from electricity’, *Environmental and Energy Policy and the Economy* **4**(1), 91–124.
- Hesterberg, T. W., Bunn, W. B., McClellan, R. O., Hamade, A. K., Long, C. M. & Valberg, P. A. (2009), ‘Critical review of the human data on short-term nitrogen dioxide (no₂) exposures: evidence for no₂ no-effect levels’, *Critical reviews in toxicology* **39**(9), 743–781.
- Hotz, V. J. & Scholz, J. K. (2003), The Earned Income Tax Credit, in R. A. Moffitt, ed., ‘Means-tested transfer programs in the United States.’, number January, University of Chicago Press, chapter 3, pp. 141–187.
- Hoynes, H. W. & Patel, A. J. (2018), ‘Effective Policy for Reducing Poverty and Inequality?’, *Journal of Human Resources* **53**(4), 859–890.
- Hsiang, S. (2016), ‘Climate econometrics’, *Annual Review of Resource Economics* **8**, 43–75.
- Hsu, A., Sheriff, G., Chakraborty, T. & Manyà, D. (2021), ‘Disproportionate exposure to urban heat island intensity across major us cities’, *Nature communications* **12**(1), 2721.
- Hua, Y., Partridge, M. & Sun, W. (2023), ‘Pollution effects of place-based policy: Evidence from china’s development-zone program’, *Journal of Regional Science* **63**(3), 703–727.
- Imhoff, M. L., Zhang, P., Wolfe, R. E. & Bounoua, L. (2010), ‘Remote sensing of the urban heat island effect across biomes in the continental usa’, *Remote sensing of environment* **114**(3), 504–513.

BIBLIOGRAPHY

- Jackson, C. K., Johnson, R. C. & Persico, C. (2016), ‘The effects of school spending on educational and economic outcomes: Evidence from school finance reforms’, *The Quarterly Journal of Economics* **131**(1), 157–218.
- Jardim, E., Long, M. C., Plotnick, R., Van Inwegen, E., Vigdor, J. & Wething, H. (2022), ‘Minimum-wage increases and low-wage employment: Evidence from seattle’, *American Economic Journal: Economic Policy* **14**(2), 263–314.
- Jardim, E., Long, M. C., Plotnick, R., Vigdor, J. & Wiles, E. (2024), ‘Local minimum wage laws, boundary discontinuity methods, and policy spillovers’, *Journal of Public Economics* **234**, 105131.
- Jin, M. & Dickinson, R. E. (2010), ‘Land surface skin temperature climatology: Benefitting from the strengths of satellite observations’, *Environmental research letters* **5**(4), 044004.
- Kaestner, R., Korenman, S. & O’Neill, J. (2003), ‘Has Welfare Reform Changed Teenage Behaviors?’, *Journal of Policy Analysis and Management* **22**(2), 225–248.
- Kahn, M. E. & Walsh, R. (2015), ‘Cities and the environment’, *Handbook of regional and urban economics* **5**, 405–465.
- Katz, L. F. & Krueger, A. B. (1992), ‘The effect of the minimum wage on the fast-food industry’, *ILR Review* **46**(1), 6–21.
- Kim, K.-H., Kabir, E. & Kabir, S. (2015), ‘A review on the human health impact of airborne particulate matter’, *Environment international* **74**, 136–143.
- Kleven, H. (2019), The eitc and the extensive margin: A reappraisal, Technical report, National Bureau of Economic Research.
- Knittel, C. R., Miller, D. L. & Sanders, N. J. (2016), ‘Caution, drivers! children present: Traffic, pollution, and infant health’, *Review of Economics and Statistics* **98**(2), 350–366.
- Koster, H. R. & Van Ommeren, J. (2019), ‘Place-based policies and the housing market’, *Review of Economics and Statistics* **101**(3), 400–414.
- Ladd, H. F. (1994), ‘Spatially targeted economic development strategies: do they work?’, *Cityscape* **1**(1), 193–218.
- Lafortune, J., Rothstein, J. & Schanzenbach, D. W. (2018), ‘School finance reform and the distribution of student achievement’, *American Economic Journal: Applied Economics* **10**(2), 1–26.

BIBLIOGRAPHY

- Leonetti, A., Peansukwech, U., Charnnarong, J., Cha'on, U., Suttiprapa, S. & Anutrakulchai, S. (2024), 'Effects of particulate matter (pm2. 5) concentration and components on mortality in chronic kidney disease patients: a nationwide spatial-temporal analysis', *Scientific Reports* **14**(1), 16810.
- Leung, J. H. (2021), 'Minimum wage and real wage inequality: Evidence from pass-through to retail prices', *Review of Economics and Statistics* **103**(4), 754–769.
- Li, C. & Managi, S. (2021), 'Contribution of on-road transportation to pm2. 5', *Scientific reports* **11**(1), 21320.
- Liang, F., Liu, F., Huang, K., Yang, X., Li, J., Xiao, Q., Chen, J., Liu, X., Cao, J., Shen, C. et al. (2020), 'Long-term exposure to fine particulate matter and cardiovascular disease in china', *Journal of the American College of Cardiology* **75**(7), 707–717.
- Liddell, T., Boser, A. S., Orofino, S., Mangin, T. & Carleton, T. (2025), 'stagg:: A data pre-processing r package for climate impacts analysis', *Environmental Modelling & Software* **183**, 106202.
- Macintyre, H. L., Heaviside, C., Cai, X. & Phalkey, R. (2021), 'The winter urban heat island: Impacts on cold-related mortality in a highly urbanized european region for present and future climate', *Environment International* **154**, 106530.
- MaCurdy, T. (2015), 'How effective is the minimum wage at supporting the poor?', *Journal of Political Economy* **123**(2), 497–545.
- Mamkhezri, J. & Khezri, M. (2024), 'Vehicle miles traveled induced demand, rebound effect, and price and income elasticities: A us spatial econometric analysis', *Transport Policy* **158**, 224–240.
- Manoli, D. & Turner, N. (2018), 'Cash-on-Hand and College Enrollment: Evidence from Population Tax Data and the Earned Income Tax Credit', *American Economic Journal: Economic Policy* **10**(2), 242–271.
- Manoli, G., Fatichi, S., Bou-Zeid, E. & Katul, G. G. (2020), 'Seasonal hysteresis of surface urban heat islands', *Proceedings of the National Academy of Sciences* **117**(13), 7082–7089.
- Manoli, G., Fatichi, S., Schl pfer, M., Yu, K., Crowther, T. W., Meili, N., Burlando, P., Katul, G. G. & Bou-Zeid, E. (2019), 'Magnitude of urban heat islands largely explained by climate and population', *Nature* **573**(7772), 55–60.
- Martin, D. & Terms, A. (2015), 'Does a high minimum wage spur low-skilled emigration?', *Economics Letters* **137**, 200–202.

BIBLIOGRAPHY

- Maxfield, M. (2013), ‘The effects of the earned income tax credit on child achievement and long-term educational attainment’, *Job market paper*. <https://msu.edu/maxfield/20131114%20Maxfield%20EITC%20Child%20Education.pdf>.
- McDuffie, E. E., Martin, R. V., Spadaro, J. V., Burnett, R., Smith, S. J., O’Rourke, P., Hammer, M. S., van Donkelaar, A., Bindle, L., Shah, V. et al. (2021), ‘Source sector and fuel contributions to ambient pm_{2.5} and attributable mortality across multiple spatial scales’, *Nature communications* **12**(1), 1–12.
- McKinnish, T. (2017), ‘Cross-state differences in the minimum wage and out-of-state commuting by low-wage workers’, *Regional Science and Urban Economics* **64**, 137–147.
- Meyer, B. D. & Rosenbaum, D. T. (2001), ‘Welfare, the earned income tax credit, and the labor supply of single mothers’, *Quarterly Journal of Economics* **116**(3), 1063–1114.
- Migration Policy Institute, . (2019), ‘U.S. Immigration Trends’. <https://www.migrationpolicy.org/programs/data-hub/us-immigration-trends>, accessed 2020-09-29.
- Miller, S. & Wherry, L. R. (2019), ‘The Long-Term Effects of Early Life Medicaid Coverage’, *Journal of Human Resources* **54**(3), 785–824.
- Minton, R. & Wheaton, B. (2022), Minimum wages and internal migration, Technical report, Harvard Working Paper Series.
- Moffitt, R. (2007), ‘Welfare reform: The US experience’, *Swedish Economic Policy Review* **14**(2), 11–47.
- Moffitt, R. A. (1983), ‘An economic model of welfare stigma’, *American Economic Review* **73**(5), 1023–1035.
- Moffitt, R. A. (2003), *Means-tested transfer programs in the United States*, University of Chicago Press Chicago:.
- Monras, J. (2019), ‘Minimum wages and spatial equilibrium: Theory and evidence’, *Journal of Labor Economics* **37**(3), 853–904.
- Moretti, E. (2024), ‘Place-based policies and geographical inequalities’, *Oxford Open Economics* **3**, i625–i633.
- Moretti, E. & Neidell, M. (2011), ‘Pollution, health, and avoidance behavior evidence from the ports of los angeles’, *Journal of human Resources* **46**(1), 154–175.
- Mullins, J. T. & White, C. (2019), ‘Temperature and mental health: Evidence from the spectrum of mental health outcomes’, *Journal of health economics* **68**, 102240.

BIBLIOGRAPHY

- Murnane, R. J. (2013), 'Us high school graduation rates: Patterns and explanations', *Journal of Economic Literature* **51**(2), 370–422.
- Neumark, D. & Kolko, J. (2010), 'Do enterprise zones create jobs? evidence from california's enterprise zone program', *Journal of Urban Economics* **68**(1), 1–19.
URL: <https://www.sciencedirect.com/science/article/pii/S0094119010000033>
- Neumark, D. & Nizalova, O. (2007), 'Minimum wage effects in the longer run', *Journal of Human resources* **42**(2), 435–452.
- Neumark, D., Schweitzer, M. & Wascher, W. (2004), 'Minimum wage effects throughout the wage distribution', *Journal of human Resources* **39**(2), 425–450.
- Neumark, D. & Simpson, H. (2015), Place-based policies, in 'Handbook of regional and urban economics', Vol. 5, Elsevier, pp. 1197–1287.
- Nichols, A. & Rothstein, J. (2016), The Earned Income Tax Credit, in 'Economics of Means-Tested Transfer Programs in the United States', Vol. I, University of Chicago Press, pp. 1–83.
- Nowak, D. J., Hirabayashi, S., Bodine, A. & Greenfield, E. (2014), 'Tree and forest effects on air quality and human health in the united states', *Environmental pollution* **193**, 119–129.
- Nowak, D., Noble, M., Sisinni, S. & Dwyer, J. (2001), 'People and trees: Assessing the us urban forest resource', *Journal of Forestry* **99**(3), 37–42.
URL: <https://doi.org/10.1093/jof/99.3.37>
- Quan, J., Chen, Y., Zhan, W., Wang, J., Voogt, J. & Wang, M. (2014), 'Multi-temporal trajectory of the urban heat island centroid in beijing, china based on a gaussian volume model', *Remote Sensing of Environment* **149**, 33–46.
- Rentschler, J. & Leonova, N. (2023), 'Global air pollution exposure and poverty', *Nature communications* **14**(1), 4432.
- Rizwan, A. M., Dennis, L. Y. & Chunho, L. (2008), 'A review on the generation, determination and mitigation of urban heat island', *Journal of environmental sciences* **20**(1), 120–128.
- Rossi-Hansberg, E., Sarte, P.-D. & Owens III, R. (2010), 'Housing externalities', *Journal of political Economy* **118**(3), 485–535.
- Schwarz, N., Lautenbach, S. & Seppelt, R. (2011), 'Exploring indicators for quantifying surface urban heat islands of european cities with modis land surface temperatures', *Remote Sensing of Environment* **115**(12), 3175–3186.

BIBLIOGRAPHY

Services, U. D. o. H. & Human (1995), ‘No Title’.

URL: <https://permanent.fdlp.gov/gpo87383/gpo87383/content.htm>

Shaefer, H. L., Edin, K., Fusaro, V. & Wu, P. (2020), ‘The decline of cash assistance and the well-being of poor households with children’, *Social Forces* **98**(3), 1000–1025.

Shapiro, J. S. & Walker, R. (2021), Where is pollution moving? environmental markets and environmental justice, in ‘AEA papers and proceedings’, Vol. 111, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 410–414.

Shirley, P. (2018), ‘The response of commuting patterns to cross-border policy differentials: Evidence from the american community survey’, *Regional Science and Urban Economics* **73**, 1–16.

Simonsohn, U., Simmons, J. P. & Nelson, L. D. (2020), ‘Specification curve analysis’, *Nature Human Behaviour* .

Tan, J., Zheng, Y., Tang, X., Guo, C., Li, L., Song, G., Zhen, X., Yuan, D., Kalkstein, A. J., Li, F. et al. (2010), ‘The urban heat island and its impact on heat waves and human health in shanghai’, *International journal of biometeorology* **54**, 75–84.

Turner, M. C., Krewski, D., Pope III, C. A., Chen, Y., Gapstur, S. M. & Thun, M. J. (2011), ‘Long-term ambient fine particulate matter air pollution and lung cancer in a large cohort of never-smokers’, *American journal of respiratory and critical care medicine* **184**(12), 1374–1381.

US Department of Health and Human Services, . (2001), ‘Status Report on Research on the Outcomes of Welfare Reform’. <http://aspe.hhs.gov/hsp/welf-ref-outcomes01/index.htm>., accessed 2020-09-29.

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

Wang, C., Tu, Y., Yu, Z. & Lu, R. (2015), ‘Pm2. 5 and cardiovascular diseases in the elderly: an overview’, *International journal of environmental research and public health* **12**(7), 8187–8197.

Wasmer, E. & Zenou, Y. (2006), ‘Equilibrium search unemployment with explicit spatial frictions’, *Labour Economics* **13**(2), 143–165.

White, C. (2017), ‘The dynamic relationship between temperature and morbidity’, *Journal of the Association of Environmental and Resource Economists* **4**(4), 1155–1198.

- Woodruff, T. J., Parker, J. D. & Schoendorf, K. C. (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.
- Wursten, J. & Reich, M. (2023), ‘Racial inequality in frictional labor markets: Evidence from minimum wages’, *Labour Economics* **82**, 102344.
- Xing, J., Hu, Z., Xia, F., Xu, J. & Zou, E. (2023), Urban forests: Environmental health values and risks, Technical report, National Bureau of Economic Research.
- Xing, Y.-F., Xu, Y.-H., Shi, M.-H. & Lian, Y.-X. (2016), ‘The impact of pm_{2.5} on the human respiratory system’, *Journal of thoracic disease* **8**(1), E69.
- Xue, T., Tong, M., Li, J., Wang, R., Guan, T., Li, J., Li, P., Liu, H., Lu, H., Li, Y. et al. (2022), ‘Estimation of stillbirths attributable to ambient fine particles in 137 countries’, *Nature communications* **13**(1), 6950.
- Zhong, N., Cao, J. & Wang, Y. (2017), ‘Traffic congestion, ambient air pollution, and health: Evidence from driving restrictions in beijing’, *Journal of the Association of Environmental and Resource Economists* **4**(3), 821–856.
- Zhou, B., Rybski, D. & Kropp, J. P. (2013), ‘On the statistics of urban heat island intensity’, *Geophysical research letters* **40**(20), 5486–5491.
- Ziliak, J. P. (2008), Filling the poverty gap, then and now, *in* ‘Frontiers of family economics’, Emerald Group Publishing Limited.