Mobility, Mixing and Urban Adaptation

Andrew Renninger

University College London
Centre for Advanced Spatial Analysis

Submitted for the degree of Doctor of Philosophy

September 2025

I, Andrew Renninger, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Cities are social reactors: they turn proximity into interactions, and interactions into recombinant ideas, yet who meets whom is uneven—and now threatened by remote work and climate stress. This thesis measures human interaction and urban structure from GPS traces, maps the mesoscale structure that shapes encounters, and tests how technological and climatological shocks distort urban mobility and mixing. First, I quantify "experienced" segregation from activity and identify broad mesoscale patterns across American cities: rings of isolation around cores and pockets of segregation, which export and import diversity asymmetrically. Race, class, and centrality are the key predictors, and these patterns endure over time. Second, I separate mixing from bridging—differentiating locations that host diversity because they are situated between communities and those that draw diversity in excess of spatial location. We then explore the ways that the pandemic and transition to remote work have durably altered mobility; constructing and monitoring a panel of data that extends for six years, we see that larger cities show greater deviations while smaller cities have snapped back to the patterns that defined them before the pandemic. Decomposing mobility into three components according to the length of the journey shows that longer trips and central places disproportionately account for the longest trips; when work recentres on the home, distances will fall but mixing could fall as well. Third, I couple mobility with weather to model behavioural adaptation to heat across multiple countries. Hot days reduce and retime activity—especially midday—while evening substitution occurs in some contexts and the effects on the poor also vary across developing and developed countries. Changes during extreme heat flatten social mixing when density would typically amplify encounters. Projections suggest smaller cities face larger relative losses, while larger cities bear greater absolute impacts without adaptation. Together, these results show how mesoscale structure, centrality, and daily rhythms govern realised contact—and how remote work and rising temperature rewire those mechanisms.

Impact statement

Within academia. The thesis advances measurement and mechanism in urban science. Methodologically, it formalises experienced segregation from activity-space exposures and mixing matrices, and distinguishes mixing from bridging to locate venues that create intergroup contact. This sharpens inference beyond aggregate indices and helps reconcile micro-preferences with meso-scale structure. The finding of "rings" and "pockets" generalises segregation from a neighbourhood property to a mesosystem property, offering a reusable template for classifying sources and sinks of diversity across cities. This thesis proposes multiple tools to quantify changes to urban dynamics if the draw of central places weakens due to remote work and extreme heat, including a mixture-of-scales approach to mobility and a null model that allows for the comparison of expected and observed social mixing–with particular relevance to central business districts. The generic operator we create to track mobility data over long time series will allow researchers to monitor changes as cities continue to evolve. These tools are portable to urban studies, network science, economic geography, and social science.

For policy and practice. The evidence supplies levers for inclusive growth. "Rings" and "pockets" provide a map for targeting bridging amenities—placing services and third places between socially distinct zones—and for protecting central places as mixing engines. This work shows that changing incentives are changing behaviours: remote work lowers the cost of staying at home and extreme heat raises the cost of going out. Remote work has enabled the suburbanisation of day-to-day activity. Further, heat acts like a tax on mobility: activity compresses away from hot hours and short, walkable trips contract first, with stronger effects in disadvantaged areas. These deformations of daily rhythms reduce mixing, redistributes spending, and curbs activity where density normally generates spillovers. For remote work, programming could generate "buzz" and draw visitors. For heat, practical responses include shade and cooling in access paths, schedule shifts for public services, and targeting cooling and transport reliability in vulnerable districts—actions that protect both human health and urban vibrancy.

Future applications. The measurement techniques in this thesis can be embedded in municipal data observatories to monitor mixing, identify candidate bridging sites, and evaluate interventions. In particular, the use of free and open data to study heat will allow for both continued research and policy evaluation. It can also inform government policies around remote work or hybrid scheduling where cross-team exposure is a production input.

UCL Research Paper Declaration Form: referencing the doctoral candidate's own work

Manuscript: Renninger, A., OĆlery, N., & Arcaute, E. (2024). *American cities are defined by isolated rings and pockets characterized by limited socio-economic mixing*. arXiv:2407.12612. arXiv link | DOI: 10.48550/arXiv.2407.12612.

1. **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2)**:**
   Not applicable—this manuscript has not yet been published.
2. **2. For a research manuscript prepared for publication but that has not yet been published:**
   (a) **What is the current title of the manuscript?**
   American cities are defined by isolated rings and pockets characterized by limited socio-economic mixing
   (b) **Has the manuscript been uploaded to a preprint server? If 'Yes', please give a link or doi:**
   Yes—arXiv:2407.12612 (DOI: 10.48550/arXiv.2407.12612)
   (c) **Where is the work intended to be published?**
   Nature Cities (accepted; in press)
   (d) **List the manuscript's authors in the intended authorship order:**
   A. Renninger; N. OĆlery; E. Arcaute
   (e) **Stage of publication:**
   Accepted (in press)
3. **3. For multi-authored work, please give a statement of contribution covering all authors:**
   Candidate's contribution: I designed the study, ran the experiments, and performed the analysis.
   Co-authors' contribution: The other authors helped design the study, and reviewed and edited the manuscript.
4. **4. In which chapter(s) of your thesis can this material be found?**
   Chapter 3 (Analytical chapter: Rings and pockets of experienced segregation)

e-Signatures confirming that the information above is accurate:

Candidate: A. Renninger **Date:** December 18, 2025

Supervisor/Senior Author signature (where appropriate): E. Arcaute
Date: December 18, 2025

UCL Research Paper Declaration Form: referencing the doctoral candidate's own work

Manuscript: Renninger, A., Neira, M., & Arcaute, E. (2024). *The role of central places in exposure segregation*. arXiv:2408.04373. arXiv link | DOI: 10.48550/arXiv.2408.04373.

1. **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2)**:**
   Not applicable—this manuscript has not yet been published.
2. **2. For a research manuscript prepared for publication but that has not yet been published:**
   (a) **What is the current title of the manuscript?**
   The role of central places in exposure segregation
   (b) **Has the manuscript been uploaded to a preprint server? If 'Yes', please give a link or doi:**
   Yes—arXiv:2408.04373 (DOI: 10.48550/arXiv.2408.04373)
   (c) **Where is the work intended to be published?**
   Undecided (journal to be determined)
   (d) **List the manuscript's authors in the intended authorship order:**
   A. Renninger; M. Neira; E. Arcaute
   (e) **Stage of publication:**
   Pre-submission (posted as preprint on arXiv)
3. **3. For multi-authored work, please give a statement of contribution covering all authors:**
   Candidate's contribution: I designed the study, ran the experiments, and performed the analysis.
   Co-authors' contribution: The other authors helped design the study, and reviewed and edited the manuscript. MN also developed the theoretical justification for our exposure segregation method.
4. **4. In which chapter(s) of your thesis can this material be found?**
   Chapter 4 (Analytical chapter: The role of bridging in mixing)

e-Signatures confirming that the information above is accurate:

Candidate: A. Renninger **Date:** December 18, 2025

Supervisor/Senior Author signature (where appropriate): E. Arcaute
Date: December 18, 2025

UCL Research Paper Declaration Form: referencing the doctoral candidate's own work

Manuscript: Renninger, A., Holubowska, O., & Blanchard, P. (2024). *Remote sensing and GPS mobility reveal heat's impact on human activity across diverse climates*. arXiv:2409.20437. arXiv link | DOI: 10.48550/arXiv.2409.20437.

1. **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2)**:**
   Not applicable—this manuscript has not yet been published.
2. **2. For a research manuscript prepared for publication but that has not yet been published:**
   (a) **What is the current title of the manuscript?**
   Remote sensing and GPS mobility reveal heat's impact on human activity across diverse climates
   (b) **Has the manuscript been uploaded to a preprint server? If 'Yes', please give a link or doi:**
   Yes—arXiv:2409.20437 (DOI: 10.48550/arXiv.2409.20437)
   (c) **Where is the work intended to be published?**
   EPJ Data Science (under review)
   (d) **List the manuscript's authors in the intended authorship order:**
   A. Renninger; O. Holubowska; P. Blanchard
   (e) **Stage of publication:**
   Under review
3. **3. For multi-authored work, please give a statement of contribution covering all authors:**
   Candidate's contribution: I designed the study, ran the experiments, and performed the analysis.
   Co-authors' contribution: The other authors helped design the study, and reviewed and edited the manuscript.
4. **4. In which chapter(s) of your thesis can this material be found?**
   Chapter 8 (Analytical chapter: Climate, heat and adaptation)

e-Signatures confirming that the information above is accurate:

Candidate: A. Renninger **Date:** December 18, 2025

Supervisor/Senior Author signature (where appropriate): No senior author.
Date: December 18, 2025

UCL Research Paper Declaration Form: referencing the doctoral candidate's own work

Manuscript: Renninger, A., & Cabrera, C. (2025). *Extreme heat reduces and reshapes urban mobility*. arXiv:2501.03978. arXiv link | DOI: 10.48550/arXiv.2501.03978.

1. **1. For a research manuscript that has already been published** (if not yet published, please skip to section 2)**:**
   Not applicable—this manuscript has not yet been published.
2. **2. For a research manuscript prepared for publication but that has not yet been published:**
   (a) **What is the current title of the manuscript?**
   Extreme heat reduces and reshapes urban mobility
   (b) **Has the manuscript been uploaded to a preprint server? If 'Yes', please give a link or doi:**
   Yes—arXiv:2501.03978 (DOI: 10.48550/arXiv.2501.03978)
   (c) **Where is the work intended to be published?**
   PNAS Nexus (under revision after peer review; revise and resubmit)
   (d) **List the manuscript's authors in the intended authorship order:**
   A. Renninger; C. Cabrera
   (e) **Stage of publication:**
   Under revision (revise & resubmit)
3. **3. For multi-authored work, please give a statement of contribution covering all authors:**
   Candidate's contribution: I designed the study, ran the experiments, and performed the analysis.
   Co-authors' contribution: The other authors helped design the study, and reviewed and edited the manuscript.
4. **4. In which chapter(s) of your thesis can this material be found?**
   Chapter 7 (Analytical chapter: Extreme heat and urban structure)

e-Signatures confirming that the information above is accurate:

Candidate: A. Renninger **Date:** December 18, 2025

Supervisor/Senior Author signature (where appropriate): C. Cabrera
Date: December 18, 2025

Papers included in this thesis

1. Renninger, A., O'Clery, N., & Arcaute, E. (2024). *American cities are defined by isolated rings and pockets characterized by limited socio-economic mixing*. arXiv:2407.12612. arXiv link | DOI: 10.48550/arXiv.2407.12612.

2. Renninger, A., Neira, M., & Arcaute, E. (2024). *The role of central places in exposure segregation*. arXiv:2408.04373. arXiv link | DOI: 10.48550/arXiv.2408.04373.

3. Renninger, A., Holubowska, O., & Blanchard, P. (2024). *Remote sensing and GPS mobility reveal heat's impact on human activity across diverse climates*. arXiv:2409.20437. arXiv link | DOI: 10.48550/arXiv.2409.20437.

4. Renninger, A., & Cabrera, C. (2025). *Extreme heat reduces and reshapes urban mobility*. arXiv:2501.03978. arXiv link | DOI: 10.48550/arXiv.2501.03978.

Contents

1 The puzzle

The world economy is "spiky" [188]: a fraction of all metropolitan areas now capture the lion's share of economic growth, innovation and even infrastructure. Half of the world's economic production occurs on just 1% of its land area [264], just 1% of cities on earth contain 21% of impervious surface area, 23% of built height, 25% of the built volume, and 27% of the population [309]. The 10 largest airports—constituting just 0.03% of all airports globally—handle 10% of all passenger traffic [10]; the 20 busiest seaports handle 50% of shipped containers annually [254]. The 20 largest cities in the world now have a combined population of 500 million, more than the European Union or the United States [309]. The spikiness of the world is on full display in the world of digital bits, rather than physical goods: 85% of technology "unicorns", companies valued at over $1 billion, huddle together in just 50 cities (see Fig. 1.1) [141]; 85% of awarded patents come from 100 cities [154]. Further, the world economy is becoming spikier as inequalities both within and between cities have risen since the turn of this century [349]. Human mobility is the mechanism driving these spikes: by compressing time and space, cities are efficient media for production, trade, and innovative recombination.

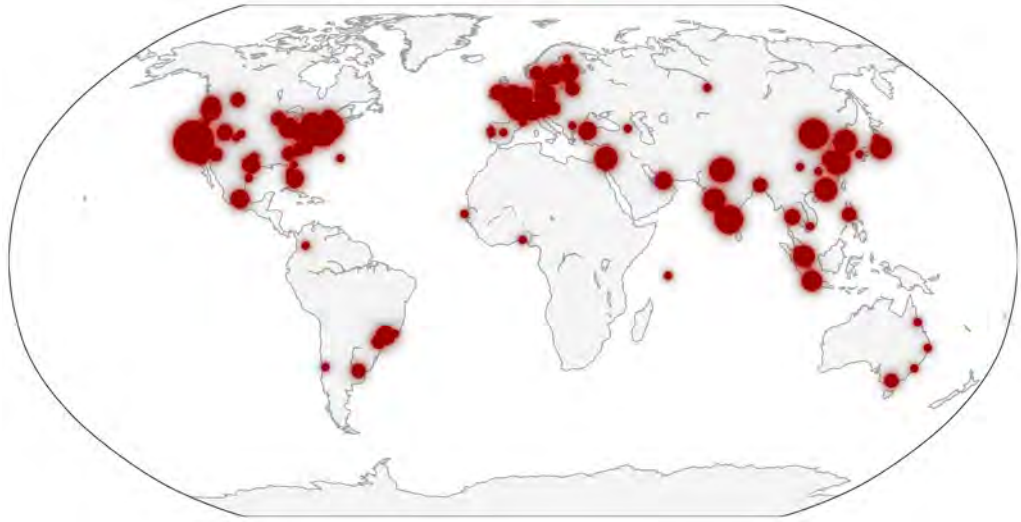

Figure 1.1. **Locations of technology "unicorns".** Although all wealth is concentrated in cities, only a few of these cities contribute to the "startup" ecosystem that has become a driver of the modern economy; all technology unicorns—companies valued more than $1 billion—reside in 100 cities, and many of the most valuable concentrate in New York City, San Francisco, London, Paris, Shenzhen, and Shanghai. Data are from Crunchbase [141].

Yet while the world is spiky and growing spikier, the spikes' locations keep shifting: the rise of San Francisco coincides with the fall of Detroit, and although absolute numbers still favour Britain's largest cities, Milton Keynes—founded in just 1967—now contributes more value per worker than Manchester or Birmingham [105]. Urban dy-

namics are at once "fast, slow, and still": firms and workers change cities quickly, cities slide up and down the rankings for population and production slowly, but the distributions that define urban systems—often close to Zipf's Law [135]—are stable [170]. Economic transformations in the past have brought about the rise and fall of cities like Detroit and Manchester, as the "perennial gale of creative destruction" [448] has wired and rewired the global economy. Megacities like Shenzhen and Shanghai are ascendant today, but Delhi and Lagos—which could each have populations of 40 million by 2050 [496, 495]—might be the future hubs of the global economy. Generally, projections indicate that the largest cities today will not be the largest cities in 2100 [233].

Having established where spikes concentrate, we now ask how they form: through everyday movement patterns that assemble (or separate) people, goods and ideas in space and time. In what follows we treat mobility as the hinge by which urban form becomes social interaction and exchange: changing how, when, and where people move changes who meets whom, who learns what.

Research identifies innovation as the engine of economic growth [459, 422, 7], and cities play a vital role in the innovation ecosystem [319]. Further, the returns to innovation accrue not just to the firms which innovate, but to other firms [76]—more so locally [293] but globally [252] as well. This in turn fosters economic dynamism and growth in distant countries [125] as innovative firms expand, sending advanced technologies around the world. In other words, productive cities create productive firms, and productive firms generate benefits for all of us. With the possibility that ideas are getting harder to find [74], ensuring that the networks that generate innovations hold together is critically important to continued economic growth and prosperity. These innovative payoffs are mediated in practice by mobility—learning, searching, and matching across space.

Herein lies the importance of understanding urban systems, dynamics, processes: they are where our prosperity will be created, more so than any deposit of precious metals or reservoir of fossil fuels [417, 425]. This means that the evolution of cities over the coming years will determine the "haves" and "have nots", as well as important technologies that benefit all of us.

Because cities' payoffs flow through encounters, mobility is the proximate phenomenon we can observe and model; it is how policy and shocks register in daily life. We saw this during the pandemic when indicators derived from various sources of mobility data proliferated—including foot traffic [207], traffic congestion [489], dinner reservations [344], office keycard swipes [250], and even a "sandwich index" of purchases at Pret a Manger stores [212].

Urban structure, the configuration and distribution of different aspects of urban life—offices, homes, amenities—in space, and human behaviour are linked. For systems of cities, we know that centres of demand coalesce into a hierarchy, wherein certain goods are provided in large urban centres—those which require large markets to remain viable—and others are distributed throughout small settlements [198, 282]. This allows for, and is the product of, the human need to access different goods at different frequencies—things we need more often must be close by, things we need less often can be far away. Although this theory models systems of cities, preliminary evidence suggests that it manifests in cities as well [444, 17]: we need localised consumption in the form of cafes and restaurants, and centralised consumption in the form of agglomerated shopping malls, high streets, and retail parks.

Cities are also the results of emergent processes that facilitate some kinds of interactions and inhibit others. Ghettos emerged to modulate how different ethnic groups access the broader city; the eponymous ghetto nuovo allowed Jewish residents of Venice to work in the city and provide critical services—but otherwise confined them to a single island

[168]. Ghettos have since become an emergent arrangement in many cities, but they always situate residents in a manner that is close but not among, allowing access without integration. Today, across the world, cities display patterns of "experienced segregation" [278], a regularity in which people from one group are less likely to interact with people from other groups in daily routine, often as a result of some of the same forces that produced ghettos. Because cities are conceived in the literature as "social reactors" [71], where gains to innovation and productivity are the results of serendipitous collisions that integrate and recombine ideas, this presents something of a puzzle: why not integrate all members of society?

The answer may lie in human biases—discomfort with different "others" [20, 360, 361]—or in cognitive limitations—our brains cannot sustain infinite relationships [167], which ramify in space. Yet literature also suggests that we may benefit when we adjust our preferences to invite discomfort and messiness into our lives [362], which in turn implies that people do not behave rationally [246, 493, 492, 248]—and that cities might not function optimally as a result. In fact, as vast and complex systems, cities present many of the challenges that force any given person to resort to "fast and frugal heuristics" [201] rather than reasoned optimisations. After strikes disrupted service on the London Underground, forcing millions of Londoners to adjust their commute, 5% of all riders never returned to their original route [266].

The possibility that our habits are metastable and can durably change with perturbations is important because technological and climatological shocks facing many cities promise to alter many aspects of urban life going forward, inviting changes to behaviour and routine. Hotter days will make it less attractive to walk or bike, possibly to leave the home at all; easier communication will make it less compelling to commute. This, in turn, will test the function of cities as social reactors, which bring people together to engineer serendipity. The "introvert economy" [447], wherein people prefer a night in with streaming television to a night out with friends, will curb the rate of collisions.

If face-to-face interaction is the engine that drives urban economies, what happens when interaction is spatially segregated, technologically attenuated, or climatically taxed? What happens when the incentives to meet in physical space fall and the returns to entertainment in digital space rise? We seek to answer each of these questions in the thesis that follows, exploring the consequences on urban systems of variations in and changes to mobility as the world equilibrates to new geographies of remote work and extreme weather. We look at the state and evolution of cities amidst a confluence of transformative pressures, some that began long before the COVID-19 pandemic and others that could persist long after.

The proliferation of large and detailed GPS mobility data—capturing daily structures of human interactions across cities and countries—now allow us to observe patterns that were previously invisible [49]. Recent advances allow us to compare interaction biases and spatial structures across diverse urban contexts at scale, track how they evolve over time, and test their persistence under major shocks such as the COVID-19 pandemic and accelerating climate extremes. This makes the core questions of this thesis, which we present below, both tractable and urgent: if we misdiagnose these dynamics, we risk pursuing policies that entrench segregation, squander the innovative potential of diverse urban populations, and exacerbate inequalities within and between cities. If we diagnose them well, we can design interventions that preserve the generative social reactor function of cities even as technologies and climatic conditions change.

Given the aforementioned technological and climatological shocks, our study is one of mobility in equilibration and transition. The goal of this thesis is to better understand the future of cities under new regimes. We ask three questions that organise the empirical

work.

- What are the broad patterns of experienced segregation across cities and regions, and to what extent are they explained by urban structure rather than human preference?
- How has the shift to remote and hybrid work reweighted mobility away from city centres, and how do city size and the strength of secondary centres moderate that reweighting?
- What are the consequences of extreme heat for the timing and geography of mobility, and for intergroup mixing?

Our empirical analysis is grouped into three parts, each mapping onto these questions: **Segregation and Mixing**, **Remote Work and Urban Structure**, and **Changing Climates, Changing Cities**. The first part explores how urban structure influences mobility and moving into a series of shocks and dynamics. A simple rule that we explore in chapters 3 and 4 is that centres mix, peripheries sort. In the second part, because central places are so important to the social mixing that happens in cities, Chapters 5 and 6 explore the consequences for cities as we redistribute greater shares of work to remote and hybrid arrangements that do not occur in central places at all but in peripheral residences—tracing a phenomenon called the "doughnut effect". In the final part, Chapters 7 and 8, we then look further into the future by exploring how climate change influences our mobility and, by extension, the nature of the "social reactor" that cities create.

Styling notes

On "we". As when presenting this research to an audience, I use the inclusive "we" to invite the reader into the argument. Each step is laid out—visually, numerically, and graphically—so it can be followed and questioned. If the "we" ever breaks—if you do not reach the same conclusions—I have not done my job.

On format. The thesis comprises six empirical chapters that have been adapted from six papers. To match journal conventions, each chapter presents results first, followed by discussion and methods. Some light repetition is intentional to keep each chapter readable on its own.

Box 1

Many early studies in human mobility were limited in spatial context and temporal coverage. While not a methodological choice, we make an important epistemological move by aggregating data from as many sources over as many periods as possible. The insights we gain from this rest on two key assumptions, which we motivate now.

"The past is a foreign country, they do things differently there." [223] The COVID-19 pandemic and the shift to remote work put in sharper relief the fact that cities are always latent and evolving—they are never removed from the past, a phenomenon known as path dependence, nor are they finished products [62]. By looking across time, we learn about the origins of urban problems and gain confidence that the patterns we identify are robust to shocks—or, if they are not, we learn about the fragility and transience of such patterns.

Since the turn of the new millennium, the defining feature of what makes a city prosperous has been turned on its head, from advantages for production to its comparative strengths for consumption—from restaurants, cafes and bars to warmth and sun—which

is in turn driving both demographic and geographic transformation [203]. The data we exploit here begin at the height of the "superstar" city, a period of growing populations of knowledge workers and rising costs of living in cities (see Fig. 1.2). In 2018, Google acquired a site for a new headquarters in New York City and began construction on new headquarters in London. In the same year, Amazon announced its decision to build a headquarters in Washington, DC, alongside another complex in New York City. The "creative class" [190] was beginning to drive up real estate prices in trendy neighbourhoods in global cities [479]. This was a period of transition from urban revival to "urban crisis" [189], with rising inequality and falling affordability alongside segregation and gentrification in many booming cities; it was also a period defined by a rich-get-richer, "the best and the rest" dynamic.

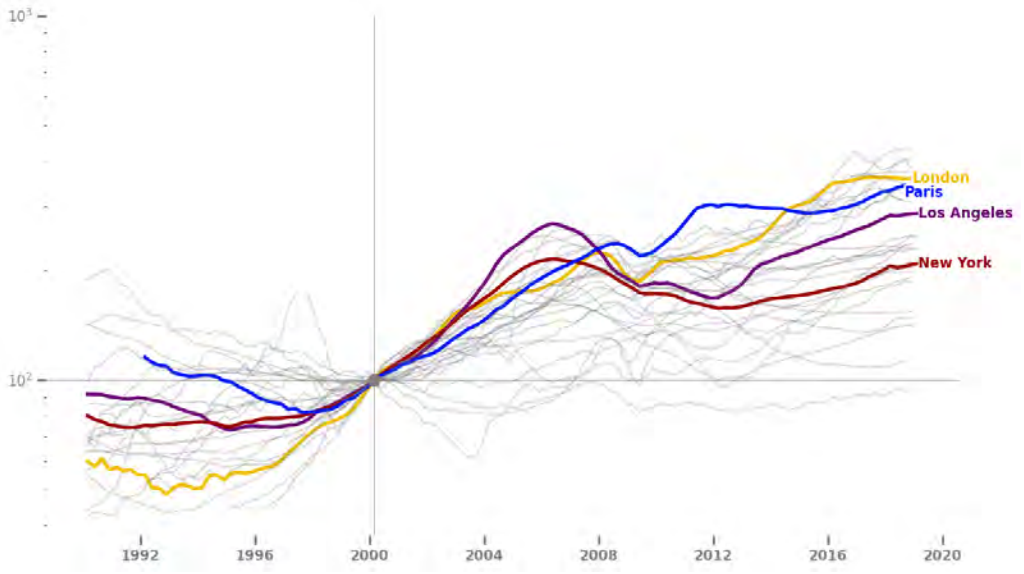

Figure 1.2. **House prices since 1990.** Beginning around the turn of the new century, cities began to grow in size, with housing becoming more expensive in many global cities; shown here, house prices are indexed to 100 in the year 2000. Rising costs of living were beginning to show in a "new urban crisis" prior to the pandemic. Data are from The Economist [478], indexed such that 2000 = 100.

The transition to hybrid or fully remote work invites new questions about what a city is for, questions that have roots in the transition from producer to consumer cities, but are accentuated by new freedoms surrounding where to live and work—if you can work from home, you might sort more on consumption and leisure amenities than if you cannot do so. As firms have moved more work to telecommunication networks, economically, cities have moved from being "space of places" to "spaces of flows" [103]: goods and services are important, but information is key and much of this information now flows through virtual links. However, this information comes in the form of both know-how and know-who, and much of it still depends on spillovers that occur in space [41]. The evidence so far is that people cannot totally divest from the city: although many have moved to suburbs or exurbs, they often remain within the same metropolitan region, because they still need to visit the city occasionally [379]—which in turn suggests that either employers or employees still consider proximity, or face-to-face interaction, important.

All of this presents a mixed narrative about how cities are different now than they were in January 2020. In light of this uncertainty, our work attempts to use long time series or comparisons from pre- and postpandemic cities to gain clarity about which patterns we see are durable, which are transient, and which appear to be the product of an interaction between local and global forces. These interactions will feature in the pages that follow

because the cities that represented the height of "superstar" urbanism—rising costs of living, growing inequalities—give critical lessons for the metastability of urban systems.

Time series allow us to calibrate our theories with the appropriate epistemic confidence. Longitudinal data reveal whether observed regularities are structural or transient, robust or fragile [149, 386]—a distinction that is critical for robust theory and for policy design. They allow us to see how shocks propagate through urban systems, how patterns persist or decay, and how feedbacks reshape trajectories [56].

"The future is already here—it's just not evenly distributed." [200] We also gain important epistemic leverage by spreading out across space. This allows us to look into possible futures, plausible counterfactuals, as well as edge and corner cases in different contexts—especially different climate zones. In the following work, we see "leading indicators" for many phenomena that we study: San Francisco has the most severe "doughnut effect" in our data, for example, and (hotter) India shows the signs of climate adaptation—displacing activity into the evening on hot days—while (cooler) Mexico does not.

This claim—that we can learn about the future by studying diverse contexts in the present—is formally called "space-for-time" substitution, and has its origins in climatology and ecology. Spatial differences can predict 75% of the temporal variation in some contexts, so long as there is strong variation across space and strong change over time [72]. While space-for-time substitutions have come under scrutiny in recent years [28], they may still add value [183]. Comparative statics do not predict a single future; rather, they identify plausible futures and allow us to identify leading indicators and emergent patterns that would be invisible from a single-city or single-country study.

One caveat: unlike ecosystems, urban systems are interlinked—and that interdependence is important to learning, the spread of ideas related to adapting to shocks. This is a feature, not a bug, of our study: that Mexico City can learn from New Delhi adds to the importance of this work. In other words, space-for-time substitutions reveal "actionable risks" in ecology [78], even when the exact predictions err, but for cities they also suggest "actionable mechanisms" by which we can mitigate those risks.

"In the science of city planning the whole city is our laboratory." [192] Here we use a variety of cities and regions to understand what drives the patterns we see in human mobility and enable better decisions for the planning and design of cities as we continue through this period of transition.

With these questions and premises in place, we begin with a review of the literature, both to situate our work and to motivate the model of cities we use. This will survey work on mobility so that we can situate our analysis within the field, but will also motivate the research and build a coherent model for how cities function. In the chapters that follow, we will consider three distinct urban transitions. Chapters 3 and 4 consider American cities before and during the pandemic—cities like New York City and San Francisco that were expensive, unequal, and divided. In these chapters, we empirically test the claim that centres mix while peripheries sort, mapping the structure of experienced segregation in Chapter 3 and probing the role of central places in Chapter 4. We then introduce methodologies for measuring urban mobility in Chapters 5 and 6, focusing on polycentricity and core-periphery dynamics, with the goal of understanding how cities change under remote work. Chapter 5 looks at the city level, employing dimensionality reduction to track mobility in 382 cities across six years. Chapter 6 looks at the neighbourhood level, using mixture models to understand the relationship between where a neighbourhood is located and how its residents distribute trips throughout the city. Finally, we test the effect of extreme heat on urban mobility in Chapters 7 and 8, with particular focus on

how this could impact cities as social reactors. Because climate stress alters incentives around work and leisure, acting as a comfort tax, Chapter 7 measures collapses in activity on hot days in Spain, and Chapter 8 extends our understanding of these behavioural responses across hotter and poorer contexts in India, Mexico and Indonesia.

Accurate measurement of mobility and intergroup contact is the difference between policy that amplifies divisions and policy that preserves the innovative potential of cities as technologies and climates shift [430, 468]. This thesis contributes in three ways.

- **Empirical.** We identify mesoscopic patterns of experienced segregation across American cities and develop models that explain these patterns. In doing so, we revise how urban structure drives experienced segregation and show that city centres play a pivotal role in social mixing. Using data with rare spatial and temporal coverage, we also show that remote work alters core-periphery dynamics, and that extreme heat suppresses mobility—both of which have consequences for social mixing.

- **Methodological.** We develop tools for monitoring cities during this period of transition so that future work can track changes under rising temperatures and remote work. We develop a technique for dimensionality reduction that allows for comparisons between cities and across time; we advance mixture models as a way to decompose mobility patterns within cities.

- **Conceptual.** We advance heat as a primary force in urban dynamics by examining its effect on mobility across both richer and poorer countries. Prior work documents consequences of extreme heat for health, productivity, time use, and travel mode; what is missing is a comprehensive account of how heat changes who moves where. We show that heat can reduce, retime, and reshape mobility, with implications for how urban economies function.

In what follows, we will measure encounters, separate structure from preference, and see what shocks do to both. If we succeed, the reader will leave with a clearer view of how to preserve mixing—and the prosperity it supports—as cities adapt to new work and a hotter world.

2 Literature

The following thesis is divided into three analytical sections. The first focuses on documenting and modelling experienced segregation—the degree to which individuals of one group are exposed to individuals of another in day-to-day activity. The second develops methods to monitor urban systems in light of changing mobility patterns that continue to evolve in present day, with particular attention given to remote work as a driver of this change. The third and final section looks at how extreme heat threatens urban mobility so that we can understand how city life might change with rising temperatures. The following review begins by motivating why we should study cities and interactions before reviewing the literature on experienced segregation, remote work, and extreme heat, mirroring the structure to come.

We espouse an organising principle: observed mixing is shaped by urban structure—the locations of homes, jobs and amenities, and the generalised cost of reaching them—and human behaviour, manifest in systematic differences in where groups go, conditioning on access. We map a set of caninical urban models to mobility in the review that follows, which we lay out in 2. Our empirical chapters operationalise this as a structure-preference decomposition of experienced segregation, and then treat remote work and extreme heat as shocks that shift incentives or raise travel costs, letting us see which patterns are mutable and which persist.

2.1 Why cities, why interactions?

A recurrent idea motivating all of our chapters is the idea that a city is a "social reactor", generating trade, innovation, even friendship by putting residents in close proximity to one another, letting goods pass and ideas spillover between them. A city represents a space of possible connections that grows exponentially with its population, and these connections give rise to certain advantages and efficiencies in production and innovation [70]. Urban scaling theory describes the relationship between a city's population and its various economic and social indicators. A variable Y relates to population, P, as

$$Y \sim P^{\beta}$$

where β represents the scaling exponent. The value of β conveys the degree to which a phenomenon grows with size, with $\beta < 1$ indicating economies of scale and $\beta > 1$ economies of agglomeration.

Research on scaling relationships in urban economies identifies three distinct categories, each with its own scaling rules [68, 31]: social aspects (wages, patents) follow superlinear scaling, infrastructure demands (roads, buildings) follow sublinear scaling, and human needs (jobs, housing units) follow linear scaling. We demonstrate this with gross metropolitan product, the total value of goods and services produced in a city, in Fig. 2.1. Simply, many of the inputs—the roads and electric cables that facilitate activity—to urban economies exhibit a sublinear relationship with city size while many

Linking theory to empirics

Understanding mobility. In what follows, we treat mobility, exposure and mixing as the product of many forces. Urban structure determines the feasible choice set and its generalised costs; incentives and constraints, as well as individual preferences, further shape realised behaviour, such that

$$\text{Mixing} \sim \underbrace{\text{form \& function}}_{\textbf{structural factors}} + \underbrace{\text{preferences, incentives \& constraints}}_{\textbf{behavioural factors}} + \underbrace{\text{transient conditions}}_{\textbf{environmental factors}} + \epsilon,$$

with ϵ representing noisy decision-making. Environmental conditions, in particular heat, act primarily by raising generalised costs and tightening feasible time windows, shifting when, where, and how people move.

Urban form and function.

- *Central place theory.* **Role:** explains why and where *shared destinations* form (centres, subcentres) and how they create arenas for encounter. **Empirical entry:** motivates core–periphery dynamics for experienced segregation (**Ch. 3–4**); provides a structural lens for decentralisation under remote work (**Ch. 5–6**); clarifies why shocks can reweight activity toward nearer centres or local hubs (**Ch. 7–8**).
- *Bid-rent theory.* **Role:** explains why activities sort with distance from centres: in equilibrium, centrality is valuable because it is access to opportunity. **Empirical entry:** interprets ring/pocket gradients in exposure (**Ch. 3–4**); frames reallocation of daytime presence and consumption when commuting weakens (**Ch. 5–6**); clarifies how travel costs shift the effective city different groups can experience (**Ch. 7–8**).

Preferences, incentives and constraints.

- *Schelling dynamics.* **Role:** illustrates how small, local preferences for like neighbours can amplify into large, global patterns without coordinated intent. **Empirical entry:** motivates interpreting deviations from structure-implied exposure as "excess" segregation (**Ch. 3–4**); mobility-coupled extensions connect sorting to destination choice [321].

Historic patterns of segregation.

- *Hypersegregation.* **Role:** describes historically produced regimes in which separation manifests across multiple dimensions (e.g., isolation, concentration, centralisation), often through durable institutional and infrastructural legacies. **Empirical entry:** disciplines interpretation of ring/pocket patterns and explains why some neighbourhoods exhibit systematically higher experienced isolation than structural factors alone would predict (**Ch. 3–4**), consistent with path-dependent constraints on access and opportunity.

Causal loops and scaling. A recurring theme in urban theory is that mobility both shapes and is shaped by cities: flows create agglomeration benefits and congestion costs, which feed back into land values and land use, reshaping the opportunity field. We summarise this co-evolution as

$$\textbf{Form}_{\mathbf{t}} \rightarrow \textbf{Choice set} \rightarrow \textbf{Mobility} \rightarrow \textbf{Agglomeration} \rightarrow \textbf{Price} \rightarrow \textbf{Form}_{\mathbf{t+1}}.$$

Within this loop, urban scaling theory illuminates the consequences: interaction intensity rises with population, making shifts in mobility and exposure plausibly consequential for aggregate outcomes.

Why think about this? This mechanism map anchors the thesis: we begin with stable imprints of urban form and function on experienced segregation (**Ch. 3–4**), then study a large economic shock that reweights activity in space (**Ch. 5–6**), and finally an environmental shock that compresses and retimes mobility (**Ch. 7–8**). The distinction matters for policy: if exposure is primarily structural, responsibility sits with the design of access and opportunities; if residual segregation persists beyond structure, levers shift toward incentives, institutions, and sorting frictions.

Box 2

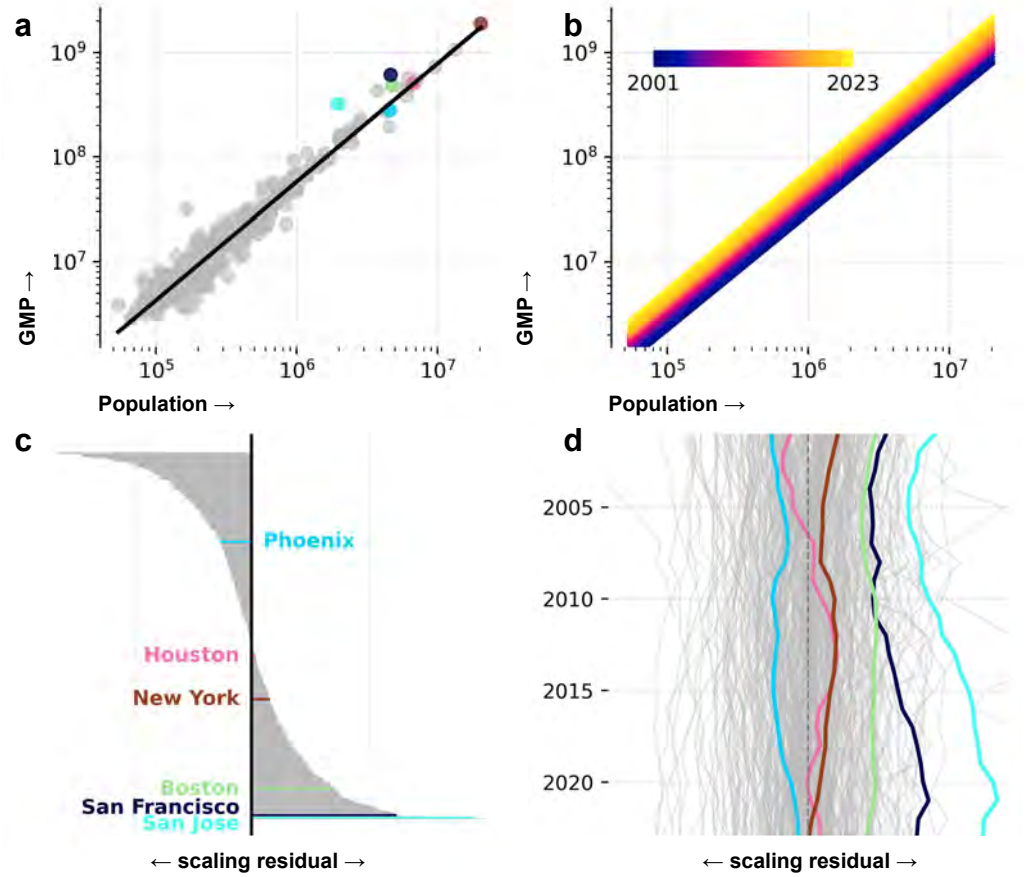

Figure 2.1. **Scaling relationships.** As US metropolitan areas grow in size, they produce more goods and services (gross metropolitan product, or GMP), but they do so at a greater rate—scaling with an exponent of $\beta \approx 1.15$, meaning that a doubling of city size corresponds to 115% rise in production, which we show in **a**. We see in **b** that this relationship is stable over time: the intercept has changed, as the economy has grown, but the slope has not. We extract the residuals in **c** to see that cities show more or less of a deviation from the scaling expectation—including San Francisco and San Jose, in Silicon Valley—which produce much more than we might expect given their populations. Important to this work, we can track these residuals across time in **d** to see that many of these residuals have changed since the pandemic, suggesting that the shock to mobility and thus connectivity is affecting the economies of these cities. Data from the US BLS [497].

outputs—the wages, patents, and goods that the city generates—show a superlinear relationship. Yet many of these regularities allow for different mechanisms. We get returns to scale in the form of more activity with less infrastructure—save for basic human needs, like residential units, which mechanically needs to scale at a 1:1 ratio with population. Agglomeration emerges from learning, matching, and sharing [171]. Sharing maps on to the sublinear growth of infrastructure, as we share infrastructure to more efficient effect as population grows. Matching and learning map on to superlinear growth in production, as ideas spread and workers move between jobs faster to create a dynamic economy.

The social reactor is predicated on the efficient transfer of goods, ideas, or people across space. Neighbourhoods with denser street networks show greater patenting rates, suggesting the pool of serendipitous interactions or the efficiency of planned interactions help generate patentable ideas [418]. Neighbourhoods where Starbucks opens a coffee shop see more new businesses than neighbourhoods where Starbucks planned to open but did not [121]. Starbucks is an example of a "third place", between home and work [343], and 24% of startup founders reported meeting to develop the ideas for their new business in one of these locations; although the germ for the idea typically comes at work, in the

same survey most founders report that shared meeting places like these are important when starting a business [140]. When employees at technology companies in the Bay Area share third places, those companies cite each other in patents [41]. A key aspect of learning is the collision of disparate ideas, and population density is associated with patents that score low on "conventionality" [65].

Recent work has shown that the nature and structure of human interaction enable urban scaling. Models calibrated with both interactions within and between cities obtain the observed β coefficients [23, 415] and infrastructure may play a role in containing and thus incubating these interactions [317]. While these are important contributions to our understanding of cities and systems of cities in the economy, they estimate the probability of interaction theoretically without grounding it empirically. They test whether the derived exponent fits the observed outcomes, but mobility data allows us to test the inputs to that exponent; in other words, this is an inverse problem where we have the data to test directly. Research explores the importance of inter-city mobility with data [255, 269], finding a strong relationship between those connections and productivity, but comparable work on intra-urban connections is lacking. Many have used variations in urban structure to proxy for intra-urban mobility [222, 145, 9, 490], finding that these proxies often correspond to other urban indicators like segregation, productivity, unemployment, and various measures of social connection. The logic is simple: barriers and distortions that increase the real or perceived distance between parts of a city act as a tax on mobility, reducing it.

From models to data, the number of contacts and calls in someone's telephone records scales according to city size, with $\beta = 1.15$, but that triadic closure—friends who are themselves friends—is invariant [443]. Increasing connectivity but constant clustering suggests that an important component of urban scaling is maintaining a collection of villages as the metropolis grows. Rates of mental illness and COVID-19, phenomena that are derived from human contact, show a noisy but positive relationship between city size and disease burden [467, 466]. Also pointing to interactions, scaling patterns are different in India, compared to the United States, and this may be due to divisions in the structure of Indian society—by caste and gender [428]. Yet evidence from France indicates that segregation does not scale with population [136]. We will see in Chapters 3 and 4 that the relationship between city size and integration is complicated by urban structure.

Because cities scale predictably, "distance to the scaling law" is a measure of urban performance when comparing cities across a range of metrics [24]. Twinning this technique with more advanced modelling—treating the residuals as the variable of interest in a second stage of regression or machine learning—could allow us to better explain variations. Many studies (see [69] for a review) check the residuals—constructing a "scale-adjusted metropolitan indicator", demonstrated in Fig. 2.1—to see which cities fall off the trend, but few attempt to model them. Implied in the scaling literature is the importance of interaction for economies of agglomeration, because we need to make and establish connections across and between cities to achieve economies of agglomeration. These models have expectations about contact, and variations in contact—manifest in mobility—should, in principle, explain deviations from the modelled expectation of production.

Models of agglomeration assume strong connectivity. Alfred Marshall argued that everyday contact in dense cities, or even industrial zones within cities, mean that "the mysteries of the trade become no mystery; but are, as it were, in the air" [297]. Silicon Valley in particular developed around a culture of intense social and spatial connectivity. These interactions were both formal and informal, planned and unplanned—university

clubs, social events, board meetings, phone calls and lunch meetings. Perhaps the most famous example of incubation through "energised crowding" [458] is the Wagon Wheel, a Mountain View bar where engineers would draw out semiconductor designs on napkins. The result was an innovation ecosystem that thrived on combination and recombination, both of ideas—as people shared knowledge—and minds—as people moved from firm to firm, often holding many jobs in a few years [434]. Today, there is evidence that these same interactions matter at the firm level [41], when employees of firms in the Bay Area share visits to cafes and restaurants, those firms cite each other in more patents relative to firms with less contact between employees.

Cities then thrive because they allow "what you know" to be enhanced by "who you know"—matching skills to needs, workers to jobs, and ideas to other ideas. If social interaction is an essential feature of cities, what determines who mingles with whom? Data from online social networks show that the probability of being friends with another member of that network decays with distance [45], suggesting that space plays a role. Cities, then, foster relationships by a combination of size N, speed V, and sprawl D—ample residents, reachable in short times, either because density reduces the distance between residents or because transportation infrastructure shrinks the effective distance between them. We show this diagrammatically in Fig. 2.2. People who are linked via transit are much more likely to be friends online, and new transit schedules that decrease travel time between contacts increase the number of calls between them—counterintuitive if we are to believe that we no longer need to physically meet our friends to maintain those relationships [89]. Instead, bursts of calls precede meetings in space [99]. Recent literature also indicates that geographic divides between neighbourhoods—like motorways or railways—reduce the likelihood that people become friends [9, 490]. Although barriers will naturally distort travel times, there also appears to be more at play, because the effect of barriers is exacerbated by differences in race and class [9].

In summary, social connections in cities, according to our best evidence, rely on proximity, accessibility, and barriers. First, the probability of a friendship between two residents decays with distance but less so when there is a transit connection [46, 89]. Second, while barriers enable sorting [294, 25], they also appear to disrupt either our ability to form or maintain relationships [9, 490]. Third, residents in larger cities have more reciprocated "friendships"—defined as those with phone calls in both directions–than those in smaller cities, with a scaling exponent of $\beta \approx 1.12$ [443]. Finally, these digital links are manifest in other actions, as data show that calls beget meetings, and more calls beget more meetings [99]—that is, telecommunication and face-to-face interaction are complements, not substitutes [89].

Furthermore, social networks often display "homophily" [307], as people with demographic commonalities—race, gender, age, class—have a greater likelihood of being friends than people without those shared attributes. Even if a person has racially diverse friends, that person is more likely to discuss "important" matters—undefined in the survey in question—with those racially similar [296]. Further, contact with other groups appears to produce anxiety and discomfort [20]. People in more ethnically diverse neighbourhoods have fewer friends, suggesting that discomfort with others might discourage tie formation [374]. Indeed, classic studies of intergroup contact show that many people have "intergroup anxiety"—discomfort with having ethnoracially different others as colleagues or friends—and prefer to avoid contact with people who are not like them along ethnoracial lines [361].

An urban economy thrives when "anyone can in principle be reached by anyone else" [68]—but what if people move around less in some cities than others? What if groups are separated or segregated within a single putative city? While current theories fail to

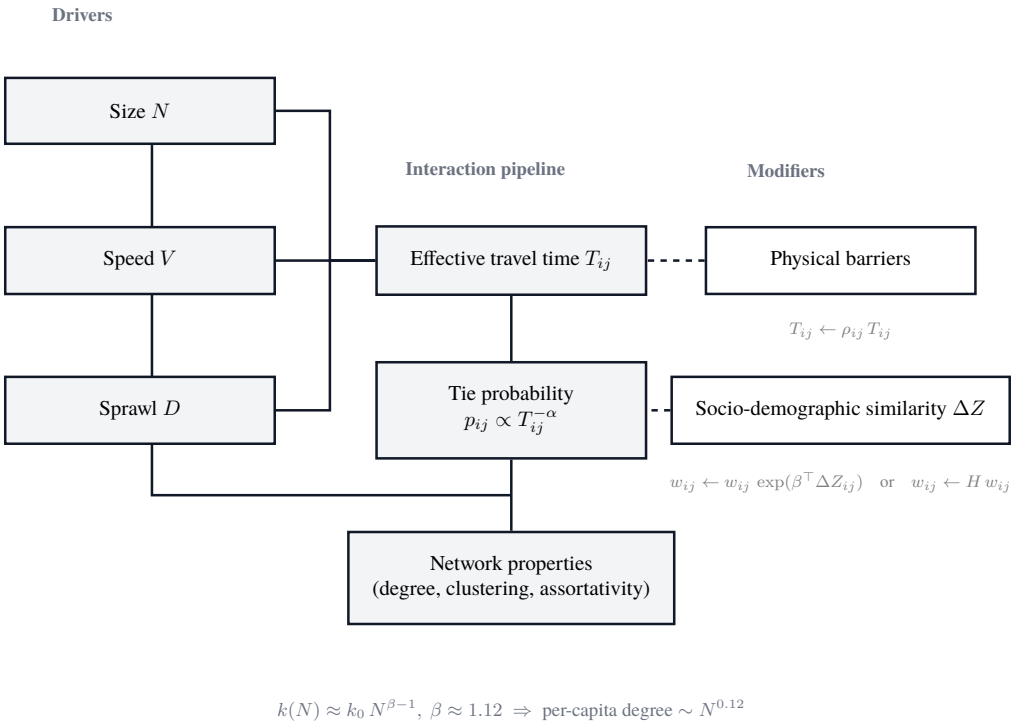

Figure 2.2. **Mechanisms linking urban structure to social ties.** "Size, speed, and sprawl" (city population N, effective network speed V, and spatial dispersion D) [373] capture how urban form and transit moderate travel times T_{ij}; tie probability decays with time as $p_{ij} \propto T_{ij}^{-\alpha}$. Modifiers act multiplicatively: barriers inflate time by $\rho_{ij} \geq 1$; socio-demographic similarity tilts weights via $\exp(\beta^\top \Delta Z_{ij})$ (or a local odds multiplier H to incorporate how similar people typically sort into neighbourhoods). A separate scaling relation $k(N) \approx k_0 N^{\beta-1}$ with $\beta \approx 1.12$ governs how average degree grows with city size; it calibrates network density, not pairwise p_{ij}.

make a distinction between inter- and intra-group interaction, mixing could be important to agglomeration—notably as a means of improving learning and matching, as well as allowing infrastructure to be shared with greater efficiency. Our research will attempt to shed light on these questions by distinguishing between mobility and integration. We will see that many aspects of mobility, including measures of who interacts with whom and how well mobility has recovered after the pandemic, scale with city size, with implications for how urban economies produce the scaling regularities that we see here.

2.2 Mobility in the city

This thesis examines how patterns of interaction emerge in cities, how these patterns are constrained by urban structure and human behaviour, and how contemporary disruptions threaten this fundamental urban function. When we discuss structure, we mean the configuration and distribution of different aspects of urban life—offices, homes, amenities—in space. Because the focus of this thesis is human mobility, and many of the decisions on where and when to travel are the product of urban structure, we begin with an understanding of how urban structure emerges.

A classic coordination game from Thomas Schelling [437] asks us to imagine a challenge: you need to meet someone whom you have never met in New York City at some time and some place, and you cannot decide where and when in advance. Where and when do you go? You might not be able to guarantee success, but you may go to the atrium at Grand Central Terminal at noon and you will likely find your partner. The sta-

tion is an example of a "focal point", a mutually-recognised, self-reinforcing belief about what others expect you to do. It is fitting that the focal point in the experiment is in a city because cities solve this coordination problem across many variants and for many actors. Of particular relevance for this work, cities provide a location—the central business district—that brings together workers and employers. As in the "pure" coordination described here, firms benefit when locating near other firms, and workers benefit when they locate where they can access those firms—ideally most or all of them. By locating all firms in the same location, at the focal point, a monocentric agglomeration reduces the mean free path length for all participants in the economy, which enables dynamism [147].

A focal point is valuable because it economises on movement, but the value of those movement economies causes land uses to segregate. Different land uses command different prices, which causes businesses and residents to sort according to "willingness to pay", producing what is called the bid-rent curve in the classic monocentric city [22, 325]. This curve answers the question, who is willing to occupy which locations, and at what price? The group that benefits most from locating in the centre will bid up prices there and those prices will deter others; annuli form as different groups sort into different locations according to how much they value proximity to the nucleus. We illustrate this in Fig. 2.3**a** and **b**. We will revisit this concentric stratification in Chapters 3 and 4 because it has consequences for travel patterns: if a suburban ring forms around an urban core, residents are at once close jobs but far from each other.

Yet cities are not monolithic entities—they are multifarious systems where different activities may thrive under different spatial configurations and scales. The locations of work and leisure activities determine much about urban life. Central place theory [122], diagrammed in Fig. 2.3**c**, posits a nested hierarchy of centres—clusters or agglomerations—with different industry mixes; some goods and services must be accessible at short distance, accessed frequently, while other markets can be farther from people because they are required only infrequently. Spatial and temporal constraints create fundamental limits on how many places we can visit in a given period. The amount of space we can traverse is limited by transport mode, and humans allocate a given "travel budget"—time spent in transit—each day; this budget is constant over time [528, 295, 435]. As a consequence, the number of unique locations a person can visit in a moving window of time is limited, and when that person adds a new location to their repertoire an old location typically falls out [18]—a new favourite restaurant crowds out an old one.

This creates an emergent spatial arrangement as businesses, populations, and possibly to a lesser extent infrastructure, adjust to meet these needs. For example, recent work suggests that the rules of central place theory hold not just between cities in larger systems of cities, but within cities, as the human need to operate across scales in characteristic ways—close and frequent, far and infrequent—extends across orders of magnitude: people will move at characteristic scales with characteristic frequencies, for example, in accordance with the same frequency–distance trade-off that central place theory formalizes [17, 444]. Daily activities draw from the immediate neighbourhood while rarer activities draw from the broader city or region. In response, restaurants cluster together to produce dining hubs throughout a city—a phenomenon that naturally occurs when regulations governing the distance between restaurants are lifted [270]. The result is a lumpy distribution of amenities, which we show for London in Fig. 2.3**d**.

This, in turn, drives changes to the population, as the suite of amenities in different neighbourhoods, in turn, influences the population distribution as people will sort *between* cities according to economic opportunities [138] but *within* cities according to amenities [195].

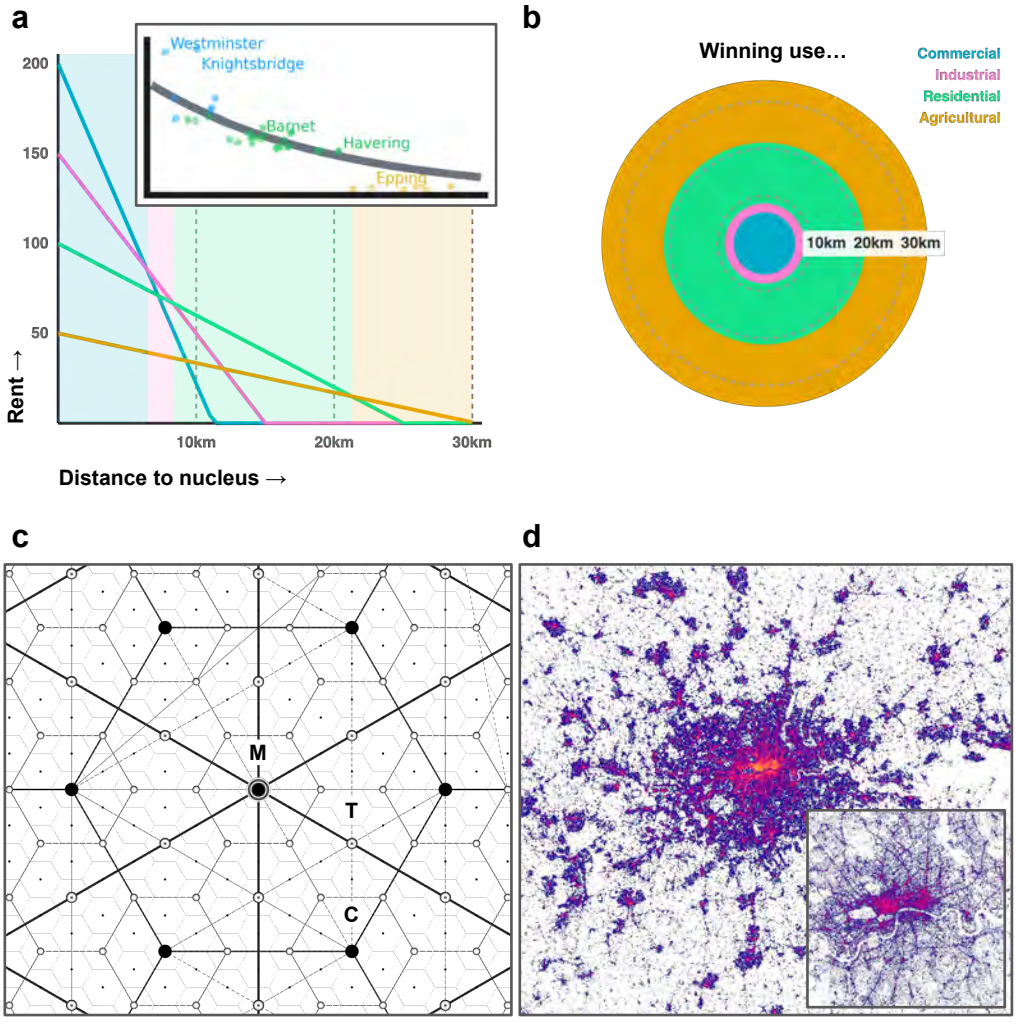

Figure 2.3. **Forces shaping cities.** In the classic bid-rent model of urban structure, shown in **a**, the benefits of a given location are a function of its proximity to the nucleus—the focal point—where the activities in a city concentrate, but these benefits will decay at different rates for different groups. We compare that to the real bid-rent curve for London. This means that the "winning" land use in a given area will be the group willing to pay most to locate in that area, creating the rings that we see in **b**. In **c** we show the predicted organisation of urban systems into a central metropole (M), peripheral cities (C), and smaller towns (T). We then map the true distribution of private businesses in Greater London in **d**, and in its centre; both global and local scales comport with the hierarchical structures of central place theory. Data are from Foursquare [193].

Because travel budgets are fixed, faster transit encourages moving away from the centre. As a result, the configuration and distribution of urban resources are never in equilibrium: over the past century, new technologies have changed the structure of cities. Steam railways separated home from work [226], electrified tramlines consolidated many businesses as faster speeds reduced the need for smaller corner stores in favour of larger grocery stores [527], and ridesharing created demand for neighbourhoods without good transit links [208]. The common theme in all of these studies is that shocks to speed and thus "reach" change the hedonic calculus for residents, and the resulting change in behaviour alters the structure of cities. Changes to the hedonic calculus—where can a person live while satisfying his or her needs—are important for understanding the shifts we see under remote work in Chapter 5. The sorting of land uses by cost has a direct consequence for mobility: the geometry of opportunities—not just metric distance—tilts flows. Classic "intervening opportunities" logic holds that flows between locations de-

pend on the count of acceptable alternatives encountered en route rather than the target's distance [469]. Modern opportunity specifications formalise this at multiple scales: the radiation family reproduces intra-urban commuting and inter-urban migration from origin and destination populations plus the intervening population [454, 455]; using population to stand in for opportunity density also predicts intra-urban trips without fitted distance kernels [524]. At the venue level, differences in trip-length distributions largely reflect how places are distributed: when distance is rescaled according to ranked intervening opportunities, place-to-place transitions follow a universal law [336]. Network speed and geometry modulate these opportunities—isochrones, not radii, determine which options are "intervening"—so embedding opportunity search on real transport networks improves predictions and recovers observed traffic patterns [387]. Street networks influence and concentrate pedestrian flows, generating "natural movement" even before attraction is added [230].

Because economising on time is a primary goal for urban residents, the locations of clusters of amenities govern how often we visit them. Recurrent visitation follows an inverse-square law in the product of travel distance and visiting frequency, linking near-daily needs to nearby options and rarer needs to farther ones—an empirical, spatio-temporal counterpart to regularities predicted via central place theory [444]. Recent work shows that the spatial "containers" shaping day-to-day movement (neighbourhood, city, region) have characteristic sizes; evidence for the sublinear expansion of activity modules with distance from home across countries further cements this nested hierarchy [17, 535]. At the individual scale, aggregate mobility patterns are a mix of distributions—in particular, the exploration of new places and the strong preferential return to a small set of familiar places. The result is a combined distribution of trip lengths with a heavy tail yet high routine; this means that in mobility high variance is coupled with high predictability [462, 463].

Finally, geometry can be separated from choice. Normalising flows by the pair distribution of locations—a continuous analogue of intervening opportunities—yields an intrinsic distance-attractiveness $\pi(r) \propto 1/r$ spanning roughly $10\,\mathrm{m}$ to $500\,\mathrm{km}$ across residential moves and day-to-day mobility; within cities, the residual kernel is piece-wise with exponents that are consistent across many centres [85]. Once flows are normalised for the number of available pairs at each separation, a simple power law collapses data across geographies, reconciling opportunity- and distance-based views [85].

Because cities are situated on a continuum between monocentric and polycentric, the structure of intervening opportunities varies across cities. This continuum spans monocentric (one core), polycentric (multiple cores), composite (one primary core, multiple subordinate cores), and an "urban villages" arrangement with distributed activities [66]. Because the distribution of opportunities matters, city structure drives mobility. The implications from this are illustrated in Fig. 2.4: with dispersed amenities, people travel less, because opportunities are close by; centralisation forces people to travel farther, but to the same place. Thus, these structures also change the probability that someone from one part of the city will interact with someone from another part. We will revisit this in Chapters 3 and 4, when we model the interactions between different groups according to the structure of the cities they inhabit. We now move from mobility mechanics to social consequences: experienced segregation—the distribution of intergroup exposure within day-to-day activities—and how it varies with structure versus preference.

On the uses, limits and biases of mobility data

What these data capture. Digital traces—be they calls and texts obtained from mobile carriers, GPS pings aggregated from mobile applications, or checkins and reviews from social media—record people at a given point in space and time [49]. Researchers represent them as trajectories, origin-destination "flow" matrices, or person-place bipartite graphs [49]; from these, we can then derive exposures—contacts between distinct groups—and visitations used throughout this thesis.

Scope and limits.

- *Behaviour.* Many studies explore regularities in mobility data—including jump lengths between discrete locations [205], daily routines manifesting across space [18], navigation choices [83]. Regularities show mobility data to be a window into human behaviour under spatial and temporal constraints. Regularities offer support for classic rules of human mobility and spatial interaction, like the importance of distance and intervening opportunities [85]. We explore this in Chapter 4.
- *Urban structure.* Scaling up and abstracting the individual in favour of the city, aggregated visits and flows reveal polycentricity and daily "dilatation" rhythms [283, 286], recover functional regions from interactions [382], and even delineate city boundaries from flows [165]. These mobility data can augment more limited studies that leverage tap-in/tap-out data from transit systems [95] by accessing a richer array of activities than those that require transit. We turn to the use of mobility data in inferring functional urban structure—as it is experienced rather than instantiated in buildings and streets—in Chapters 5 and 6.
- *Epidemic and economic signals.* During COVID-19, because viral spread is contingent on face-to-face contact, early uses of mobility data—from global flights to local trips [120]—built epidemic forecasts. Mobility data also provide a window into the social and economic consequences of the pandemic, including retail activity [114], restaurant patronage [315], experienced segregation [522], and the vulnerability of one business to the closure of another [523]. We consider shocks from the pandemic in Chapters 3, 4 and 5.
- *Inequality and exposure.* Activity-space exposure measures shift focus from where people sleep to whom they meet and where [38]. Studies leveraging these data show that spatial proximity shapes social connectivity [89] and that classic results on network advantage apply [211]. Recent work shows exposure segregation can be higher in larger cities [335]; social mixing varies systematically across amenities and streets [185]. We build on this growing body of literature to quantify what is called "experienced segregation" throughout this thesis, but notably in Chapters 3 and 4.

Data and design. Research distinguishes between *flows* and *trajectories*. Origin-destination flows quantify the spatial interaction between one location and another, often a home and some leisure, consumption or work location. From trajectories, researchers can derive "stops" or "visits" [35] and, by looking at who was in the same location at the same time, construct a list of "exposures" that capture interpersonal interaction. Research also often chooses a given *scale*, looking at human behaviours or urban patterns, the citizen or the city. In order to understand experienced segregation, we also choose *aggregation* variables—socio-economic strata or ethnicities, venues or neighbourhoods, hours or days or months—to produce mixing matrices, where each cell counts, for example, the number of contacts between groups in a given month.

Bias and best practice. Mobility data come with a variety of concerns regarding bias [197], which we mitigate with a variety of strategies in the following thesis. (i) *Representativeness* depends on who owns mobile phones and who opts into tracking from mobile applications [509]. We address this by ensuring that the number of devices across neighbourhoods corresponds to administrative data on population and does not systematically vary according to the racial and socio-economic composition of the neighbourhood. (ii) Samples also suffer from *panel drift*: which users are included in the sample change. When looking across long time horizons, an integral part of this study, we look for discontinuities that signal compositional changes. (iii) *Measurement* and semantics are not infallible: for example, buildings can distort GPS signals and put a device in a spurious location [504]. Our results do not depend on granular location, and we prefer coarser areal estimations when possible. (iv) Finally, *privacy* is an issue: location traces contain sensitive information that can identify an individual [111]; this makes aggregations that consider populations rather than devices important.

Box 3

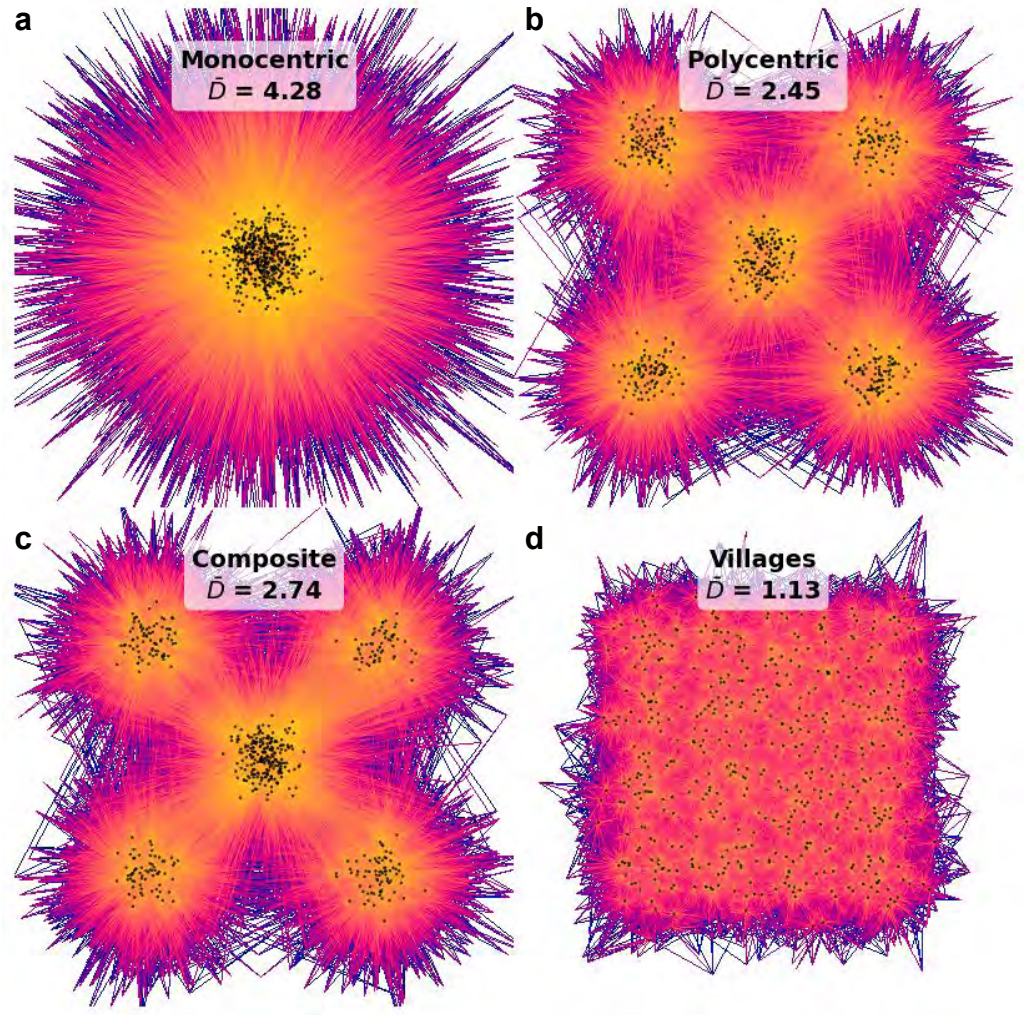

Figure 2.4. **Intervening opportunity and mobility.** Illustrative demonstration of how urban structure shapes movement patterns under the intervening opportunities model, where $\bar{D}$ represents a stylised average distance travelled. **a** In the monocentric city, central amenities draw visits from peripheral residents ($\bar{D} = 4.28$). **b** The polycentric arrangement distributes trips across 5 activity centers ($\bar{D} = 2.45$). **c** The composite city combines features of both: a central district attracts longer trips while peripheral clusters serve shorter needs ($\bar{D} = 2.74$). **d** Urban villages achieve the shortest distances ($\bar{D} = 1.13$) through local trip patterns.

2.3 Experienced segregation and urban mixing

In a classic simulation from Thomas Schelling, subtle differences in how much people prefer to be surrounded by others like them can drive large changes in sorting [438]. As tolerance for others falls, the population of a city sorts into homogeneous and contiguous zones with demographically similar residents. In this model, "there is no single person who intends that there be segregation" [439]: instead, individual preferences give rise to emergent patterns, which we demonstrate in Fig. 2.5**a**. Agents are endowed with a tolerance τ that governs how many "different" neighbours they are comfortable with, as well as a group assignment, yellow or blue; they are then allowed to move over turns until they satisfy τ. Even at moderate levels of tolerance, sorting emerges. The built environment can enable this: in the US, highways have enabled sorting by allowing large quantities of white residents to move to the suburbs while keeping access to jobs in the city centre via these conduits [59, 294], and policies that prevented minorities from moving

to certain neighbourhoods reinforced this sorting [424].

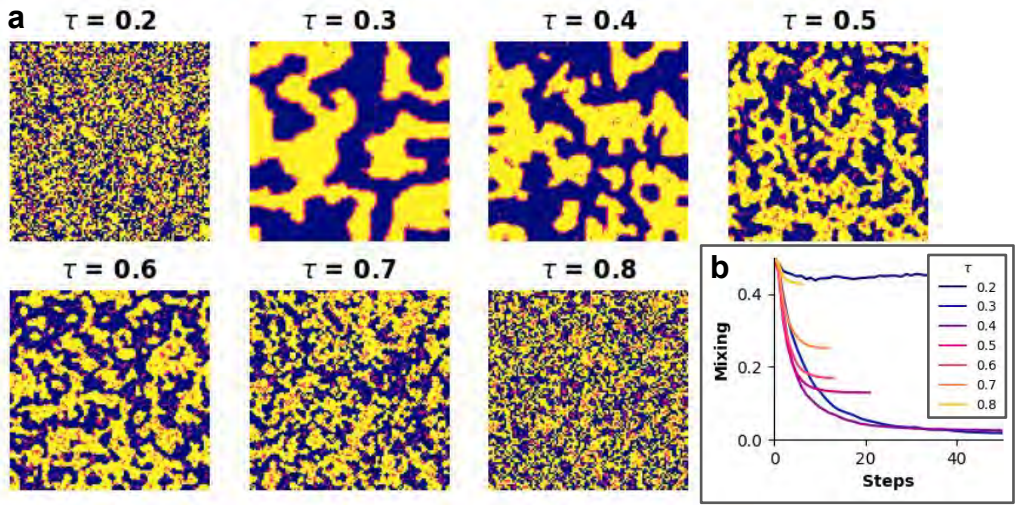

Figure 2.5. **Emergent segregation.** Illustrative model of segregation adapted from Schelling [438]. Agents of one of two groups are endowed with a tolerance τ, which determines the number of "others" with whom they are comfortable living near. At each turn they are allowed to move if the observed neighbours do not meet this tolerance. The model shows that even at moderate or high tolerances, there is still observable sorting, in **a**, in this gridded city. The time series of mixing as residents move, in **b**, shows that very low tolerance is unstable but moderate/high will result in stable configurations given a long enough time horizon.

A large body of research parses segregation into five dimensions—unevenness, isolation, clustering, concentration, and centralisation—arguing that some cities exhibit hypersegregation along multiple dimensions simultaneously [301]. In this conception, segregation concentrates minorities in contiguous spaces, isolating them from the city as a whole—often near the centre of the city. Hypersegregation concentrates poverty and severs opportunities for social and spatial mobility, producing durable inequality and "disadvantage" [512]. This canonical view frames why "exposure" matters: segregated communities depress contact with other groups and, in doing so, with new opportunities. The durable impacts of both hypersegregation and the related phenomenon of "white flight" reappear in Chapter 3.

Much of the literature on segregation treats segregation and inequality as mechanically connected, because rising inequality increases the clumping of incomes in space [508, 384]. Further research suggests that inequality enables segregation by allowing wealthier groups to bid for properties in expensive neighbourhoods to the exclusion of others [348]. Yet the degree to which segregation is good or bad is contested: many studies find inequality drives segregation (not vice versa), whereas others show segregation reduces wealth creation [274, 235, 25] and economic mobility [117, 113].

A recent turn in the literature on segregation is toward "activity space" segregation [431, 88, 96]. There are many different approaches to quantifying segregation in mobility data, which we show in Box 4, but the unifying theme in monitoring not just where people live but to which locations people go over the course of the day and looking at the diversity of people who share those locations with them. These data support much of the literature on hypersegregation, finding that communities with high crime and high poverty often have limited permeability: relative to other parts of the city, there are fewer trips out of these areas by residents, and fewer trips into them by visitors [271]. Cities with more connections between neighbourhoods have less crime [430].

Yet people who live in the same neighbourhood can have very different routines, sharing just $\sim 15\%$ of visited locations in a given month in one estimate [87]. This

means that living nearby does not guarantee shared daily experiences, and residents with higher levels of income or education will often move to different parts of the city during the day, and white and nonwhite residents often do not share the same hubs, converging on different parts of the city [102]. With opportunities to sort in local neighbourhoods, focal points—central location, downtowns—facilitate social mixing, a phenomenon that we explore in Chapter 5: in one study of how migrants interact with locals, $\sim 60\%$ of all intergroup interactions occurred downtown [420].

At large scales, mobility data reveal "experienced" isolation, a dynamic measure of interactions in daily experience between different groups compared with the static measure of segregation using residence; these measures are strongly correlated, but we are exposed to more diversity in daily life than administrative data capture [38]. Yet our activity spaces—the ground we cover through the day—reveal our race and class [96]. Segregation in online interaction networks and offline transaction networks is correlated as well, and both online and offline segregation is highest amongst the highest and lowest socio-economic strata [166]. Isolation correlates with social status but varies by time and day [38, 519], and tends to map onto neighbourhood and amenity structure [3]. As above, with growing evidence that cities are arranged to meet our mobility needs, these findings suggest that mixing by race or class is in part governed by the structure of the city.

Residential segregation is a key driver of its experiential counterpart: across many countries and contexts, the two measures are correlated [278, 38, 352]. Generally, more residentially segregated cities have more experientially isolated neighbourhoods [505]. Because many people move away from home during the day, many studies document reductions in experienced segregation during the day [38, 324, 453, 2]. Yet this appears to be driven by commutes in some contexts, as many sort into homogeneous groups for leisure activities [475].

Across multiple contexts, the middle class appear to be a bridge between lower and upper classes, both in social networks [166] and in spatial networks [166, 330]. For example, in Brazil, rich and poor are unlikely to share a common route to work, but the middle class will often make journeys and contact both [330].

Experienced segregation is typically lowest during the day [40], and a plausible mechanism is simply "population dilation" [283]—central cities absorbing residents from peripheral satellites during peak business hours. For example, the most central arrondissements in Paris grow by 150% during the day, and during the day experienced segregation is at its lowest [268]. Yet population dilation is governed by the structure of the city: the central business districts of larger, polycentric cities receive fewer visitors, as a share of the total population, than smaller, monocentric cities [283]. Mechanically, as cities grow, congestion leads to fragmentation [287], so polycentricity is the norm in large cities. By extension, larger cities should have higher rates of experienced segregation as fewer residents travel to a single focal point each day, and indeed data in the US show this [335].

Yet the structure of cities cannot explain all the patterns we see in the data. Both spatial and social frictions combine to determine the clientele of restaurants or other businesses: not only do transit times matter when measuring an individual's propensity to visit a restaurant, the clientele of that restaurant matters too [150]. Mixing patterns at points of interest in 11 cities reveal that economically integrated and segregated places coexist metres away, so preferences beyond transit must be at play [321]. Simulations that use the gravity equation to model interactions between areas do not predict observed data, suggesting that people are not simply interacting with others nearby [321, 166]. We will revisit the role of structure in Chapters 4 and 5, what it can and cannot predict.

People are more likely to visit venues popular within their own socio-economic class

Measuring segregation and experienced segregation

Approach. Empirically, "segregation" asks how far realised contacts deviate from the opportunities implied by composition. Four traditions recur: (i) *spatial-analytic* measures of evenness across areal units [385], (ii) *information-theoretic* measures of relative predictability [450, 363], (iii) *network-scientific* measures based on assortativity and spectral measures [332, 173], and, the focus of this thesis, (iv) *activity-space* exposure from mobility traces [39, 321].

Notation. Imagine we have groups $g, h \in G$, which represent socio-economic strata, ordered low→high. We have individuals i and amenities α. City composition π_g (share of people/activity in g). Each individual is exposed to different groups throughout the day, creating exposure shares $e_i^{(h)}$ representing the fraction of i's co-presence with group h; exposure to one's own group will be $e_i^{(g)}$. People will mingle at amenities, or venues, which will have shares $\tau_{g\alpha}$—the fraction of time at α contributed by g. These produce a "mixing matrix" M with entries M_{gh}, where each row and column represents a class and each cell gives the share of g's contacts with h, such that row g sums to 1.

Formulas.
(1) Experienced isolation. Define group means $E_g = \frac{1}{|g|}\sum_{i\in g} e_i^{(g)}$. The contrast between two groups g, h at the city level is

$$I^{\exp}(g,h) \;=\; E_g - E_h,$$

a time-and-place weighted analogue of residential isolation computed over actually visited microplaces and times [39]. At the individual level we use $e_i^{(g)}$ for $i \in g$; if a contrast is needed, $I_i = e_i^{(g)} - \bar{e}_{h^*}^{(g)}$ for a chosen comparison group h^*, where $\bar{e}_{h^*}^{(g)}$ is the mean exposure of comparison group h^* to group g.

(2) Amenity and individual segregation. With income quartiles $q \in \{1,2,3,4\}$ and venue shares $\tau_{q\alpha}$,

$$S_\alpha \;=\; \frac{2}{3}\sum_{q=1}^{4}\Big|\tau_{q\alpha} - \tfrac{1}{4}\Big|,$$

which rescales the L^1 distance from uniform to $[0,1]$ for four groups; analogously for an individual's exposure profile ρ_{iq},

$$S_i \;=\; \frac{2}{3}\sum_{q=1}^{4}\Big|\rho_{iq} - \tfrac{1}{4}\Big|.$$

These capture uneven time allocation across socio-economic classes at venues and for people [321].

(3) Evenness and richness. For a composition $\pi = (\pi_g)_{g\in G}$,

$$H(\pi) \;=\; -\sum_{g\in G}\pi_g \,\log \pi_g, \qquad J(\pi) \;=\; \frac{H(\pi)}{\log|G|},$$

where H is Shannon entropy [450] and J is evenness [375], normalising by $\log|G|$ so that $J \in [0,1]$. Richness is simply $|G|$, the number of groups.

Reading mixing matrices. We employ mixing matrices in what follows; these aggregate behaviours to patterns of interaction between socio-economic strata, with M_{gh} ordered by class (rows g low→high). Row g shows whom individuals from class g meet. A row identical to the city composition π indicates no sorting by g. The upper and lower triangles convey important information: mass in the *upper* triangle ($g < h$) signals *upward* bias (lower-income meeting higher-income more than π), mass in the *lower* triangle ($g > h$) signals *downward* bias; diagonal mass M_{gg} is own-group exposure. Deviations can be read row-wise as $\Delta_{gh} = M_{gh} - \pi_h$.

Related measures. We also briefly touch on measures of correlation: a summary of exposure segregation that compares one's socio-economic status to the mean status of those encountered [335]; this is connected to assortativity computed on the exposure network [332].

Box 4

[231], but there tends to be a bias: when individuals do deviate from assortative mixing, they prefer points of interest with higher status patronage—relative to their own—than lower. This preference for mixing upward rather than downward is present in other works [166, 81]. In many countries, online social networks are more assortative than spatial interactions [166] and in these networks users can control who they follow, this hints at the role of preference. Because these networks show higher assortativity than offline networks for the wealthiest and poorest groups [166], an important integrative force in cities might be the ability of the poorest to visit at least some places that the wealthiest visit, but experienced segregation persists.

Socio-demographic communities in a city will often form a "city-within-a-city". In the purist display of these preferences, the wealthiest residents of London produce "cloud spaces" or "flowing enclaves"—travelling between private clubs in private cars, for example—to avoid contact with the broader public [42]; the elite in London can spend as much as 80% of the day in exclusive zones. Yet the phenomenon is not just practiced by the elite: in Northern Ireland, Protestants and Catholics use different—ostensibly "neutral"—streets and parks [148, 4]; in Israel, even Muslim workers who commute to Jewish neighbourhoods spend $\sim 80\%$ of time away from work in their home neighbourhoods [446]. When a new rail line opened in Jerusalem, mixing between Jews and Muslims rose at stations but not throughout the city [419].

Short trips drive segregation: longer commutes are linked with more diverse friends [81] and longer trips are still assortative but are associated with more diverse interactions [166]. Because shorter trips are frequent and longer trips are infrequent—as central place theory would predict—these homophilious interactions with those nearby, supported by residentially homophilious neighbourhoods, drive more of the assortativity in spatial interaction networks than do interactions with those farther away [166].

Yet, as discussed, there is also evidence that visits are a function of intervening opportunities and that the configuration of amenities can drive our choices [336]. The distribution of amenities, clustering and defining neighbourhoods, may mean mixing is the product of both spatial structure and social bias. Indeed, research from the US shows that cities with clusters of amenities situated between rich and poor communities show greater socio-economic mixing than cities with amenities that are nested within socio-economically homogeneous areas [335]. Further, cities that offer more amenities closer to residential populations, approaching what is called the "15-minute city", see higher experienced segregation [3]. In Chapter 4 and 5, we will revisit this while considering a trade-off: dispersion decreases distance but increases isolation—a direct structure–mixing trade-off.

The largest studies explaining isolation in the US primarily use regressions [38], by extending existing mobility paradigms [322], but findings hint that the structure of the city plays a role [335]. Travel distance is inversely proportional to experienced isolation [321]; assortativity is lower for individuals with longer commutes, relative to those with shorter ones—although clustering coefficient, friends who are friends with each other is also lower [81]. Given other evidence that clustering coefficient is invariant to city size [443], this both suggests underlying heterogeneity and the possibility that some people with less triadic closure link communities in larger cities that is otherwise collections of tighter villages. Beyond these associations, the first two analytical chapters of this thesis attempt to understand what about the structure of cities drives experienced isolation, drawing on central place theory.

With greater knowledge of these drivers, we might be able to adjust small aspects of the mobility calculus to big effect. Minor changes to the spatial interaction network—rerouting just 5% of trips—can create an egalitarian regime where all neighbourhoods receive an

equal proportion of commercial spending without increasing travel distance, and this intervention also improves mixing between neighbourhoods [284]. This would correspond to socio-economic or ethnic mixing in a segregated city, but evidence for how this could work is mixed: some studies find that out-of-neighbourhood activities provide most of the observed intergroup contact [513], but others show that many people remain in areas dominated by like groups even when they leave home [452, 446].

Further, cities present a complex mix of what are called "affordances"—the set of possible actions an environment offers an individual [199]. In the presence of such a vast array of possibilities—in the case of urban environments, places to eat, sleep, work, recreate—people will resort to simple rules of thumb, called heuristics [201]. In accordance with this literature, Stanley Milgram postulated a theory of "urban overload" [312], from which we seek reprieve with cognitive economy—selective attention—and social norms, such as limited courtesies: a person in a town, for example, is much more likely to let you use his or her phone than a person in a city; the resident of a big city is much more likely to walk past someone sleeping on the street than a resident from a small town [312]. Norms around how and when to help others are part of a suite of adaptations that help urbanites manage the cognitive and social demands of life in a city.

Given our innate preferences for connecting with others like us, it is plausible that norms, heuristics and biases interact and combine to produce the patterns we see in mobility data. The data indeed suggest as much, given that the distribution of amenities cannot explain all of the variation in visitation between points of interest [321, 150]: restaurants on the same street can have very different patrons.

In summary, cities are generative systems that match skills, jobs, and ideas. People who live in cities are subject to an array of cognitive biases, psychosocial limits as well as spatial and temporal constraints; further, the innovative communities in cities might require modularity and exclusivity. This indicates that the segregated, homophilous patterns we see in urban social life are—at least to some extent—the result of an optimisation. Yet the evidence that perturbations can so often encourage individuals to durably change routines, habits and behaviours suggests that there are still shifts to be made that allow more groups to flourish. Chapters 3 and 4 quantify how much of experienced segregation is recoverable from the structure of opportunities and networks compared to preferences.

2.4 Urban structure and remote work

Before 2020, tight housing markets in "superstar cities" [218] reflected a long shift from producer to consumer urbanism [203]: amenity-rich cores commanded rising premiums and drew high-income households, reinforcing spatial sorting even within metros, sorting that we will see in Chapter 4 manifests in both housing and activity spaces. Classic evidence on "consumer cities" documents faster growth in high-amenity places and downtown amenity clustering; complementary work explains how price dynamics in "superstar" metros amplified inequality in access to those amenities [138]. Data from this period show that larger cities, naturally fragmented due to sprawl and congestion, were and still are often collections of differentiated venues that sort visitors by socio-economic status [335]. The pandemic was an exogenous shock to that equilibrium, not only depressing footfall but reshaping who meets whom. The income diversity of urban encounters fell by 15–30% and remained depressed through 2021 [522] even after aggregate mobility recovered, indicating the pandemic durably changed preferences—either for the kinds of venues that promote intergroup interaction, or for mixing itself. These findings motivate our empirical strategy later in the thesis: we will track place diversity and venue choice

across cities and neighbourhoods, testing whether the pandemic's mixing deficit persists in amenity spaces that previously served as bridges between classes.

Remote and hybrid work have durably altered urban economies. Across surveys, rates of working from home (WFH) have stabilised: the average employee in the US took 28% of working days from home in 2023, four times the number of days worked from home in 2019 [53]. Results are comparable across much of Europe and other high-income countries, where WFH receded from pandemic highs but settled into a new plateau higher than prepandemic levels [15]. Feasibility studies conducted prepandemic show that about 37% of US jobs can be performed entirely at home, helping explain the persistence of hybrid arrangements [162].

In land and housing markets, the "doughnut effect"—the reallocation of people, firms, and demand from dense urban centres to lower-density suburbs and exurbs—has been documented first in US metros and then globally, with magnitude tied to local WFH rates [379]. A parallel adaptation is the rise of supercommuting—longer trips taken fewer days per week, in a nod to central place theory—which was flagged a decade ago and has been enabled by hybrid schedules [323, 381]. The share of commutes longer than 120km in the US has risen 32% since 2020 [73].

The evidence for whether or not remote work trades productivity for comfort is mixed: early experimental evidence shows that call-center agents became about 13% more productive and quit at lower rates when assigned to WFH [75]; more recent syntheses across industries suggest that fully remote arrangements are, on average, roughly 10% less productive than fully in-person work—primarily due to communication and coordination frictions—while hybrid arrangements show little to no penalty [53]. When Microsoft went fully remote during the pandemic, internal communications showed fewer bridging ties across groups and a shift from synchronous to asynchronous communication; the rise in virtual meetings did not compensate for the loss of in-person interactions [525]. Working from home saves time that would typically be spent commuting: across 27 countries, workers save 72 minutes per WFH day, reallocating $\sim 40\%$ of that to the job but also $\sim 11\%$ to caregiving—one reason why WFH is more important to women than men [14]. Yet early work suggests that because of lower visibility, remote workers are promoted at lower rates [75]. These countervailing forces mean that cities are still equilibrating as the largest employers calibrate WFH policies. In Fig. 2.6 we show trends in both prevalence and preference: shares of remote work across large cities have fallen from pandemic highs and stabilised below the peak, while employees systematically desire more days remote than employers plan [15].

The result of this push and pull will determine urban structure in consumer cities in particular. With fewer daily commuters, activity and spending anchor around homes, not offices, with measurable spatial reallocation in people and spending [379, 19].

Telecommunications tend to complement rather than substitute for face-to-face contact: mobile communication intensity rises superlinearly with city size, link formation decays with distance, and dyads that call frequently also meet more frequently—evidence that digital tools often scaffold, rather than replace, co-located networks [443, 89, 99]. Recent evidence from Silicon Valley indicates that a 1% drop in face-to-face contact between firms reduces patent citations between them by 0.25% [41]. In light of these facts, this thesis Chapter 5 will quantify how hybrid work restructures encounters in business districts, and whether the shift toward irregular, long-distance and low-frequency commuting further polarises the geography of everyday life—including in how both retail business and spending are distributed in the new "cloud" city. Although lacking an explicit temporal dimension, Chapter 6 will build on this by introducing a method to monitor urban structure—structure that will change if remote work becomes entrenched in

professional life.

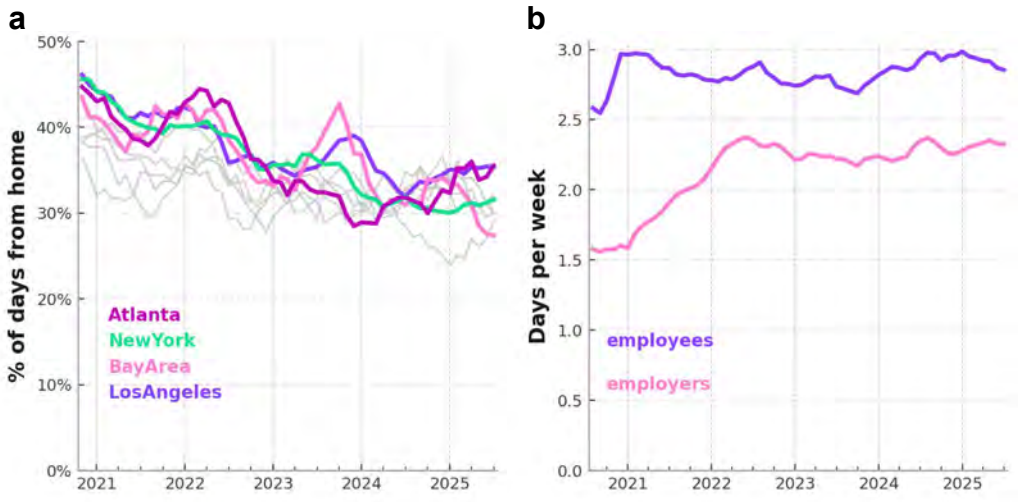

Figure 2.6. **Trends in remote work.** Practices surrounding working from home are still evolving as employers and employees negotiate policies. In **a** we see that since the pandemic, the share of days worked from home has fallen—even in technology hubs like the Bay Area. In **b** we show the preferences of employers and employees, respectively, for remote days; there is a substantive difference, bosses wanting workers back in the office and workers preferring work-from-home. The result of this negotiation will influence the structure of cities going forward. Data are from WFH Research [52].

2.5 Climate stress and city life

The next structural pressure on urban mobility is climatic. In coming years and decades, modelling suggests risk from floods [221], hurricanes [259], and fires [1] will rise. Although climate poses a number of threats to urban life, extreme heat is the most ubiquitous channel. Cities suffer from what is called the urban heat island effect, wherein temperature in cities is higher than in surrounding rural areas because of anthropogenic waste heat and absorptive surfaces such as asphalt and concrete [432]. This effect means that cities could see temperatures rise by $\sim$2°C more than surrounding areas in the coming decades [181].

Micro- and macro-evidence show that higher temperatures reduce hours worked in exposed industries [210], depress on-the-job productivity [460], and reallocate time from outdoor to indoor activities [210], with cognitive performance losses [267], which, for children, accumulate into reductions in human capital [351]. In both developed and developing countries, heatwaves are associated with measurable drags on productivity, with productivity falling by as much as 2–4% per °C on hot days [460] and with aggregate losses evident at the national level [92, 157].

At the system edge, transport infrastructure itself degrades under heat (asphalt rutting, rail buckling), increasing generalised travel costs and disruptions [164, 376, 126]. For residents, heat measurably suppresses active travel—including bikeshare usage across diverse cities [225, 61, 257]—and shifts time use toward the home [210, 292], while transit riders face disproportionate heat exposure during access and waiting [172, 311].

The net effect of heat on transport is mixed, because warm weather encourages active travel while hot weather suppresses it—for example, in New York City bikeshare usage increases up to $\sim$ 28°C and decreases beyond that mark [225]. Other evidence suggests that transit trips substitute for cycling on hot days [507], which could put pressure on public transit systems if heat begins to generally raise demand for them.

These direct and indirect frictions constitute a de facto tax on mobility. Combined with macroeconomic implications already visible in productivity losses during hot spells, the results could be profound for cities. Nevertheless, evidence suggests that access to air conditioning attenuates many of the effects of extreme heat [51, 351]. Remote work features here as well: authoritative assessments list teleworking and schedule shifts among feasible urban adaptation options that reduce exposure during heat extremes [44]. Although these mitigation strategies are individualistic, collectivistic approaches are also important and effective: cities can reduce ambient temperatures by $\sim$1°C via urban trees and by up to $\sim$2°C via reflective "cool" surfaces [181, 536].

The final empirical chapters of this thesis leverage these insights: we treat heat waves as a shock to urban mobility, estimating how rising heat shifts mode choices and destination portfolios, and run simulations to project how this trend may worsen without adequate adaptation. We also revisit social mixing. Climate shocks can reorganise encounters: for example, wildfire evacuations produce transient dispersion followed by re-concentration along socioeconomic lines, highlighting how different shocks (leave-home rather than stay-home during COVID) have opposite effects on experienced segregation [327]. Yet if heat represents a tax on mobility, its effects will hew closer to stay-home than leave-home. By raising generalised travel costs and compressing activity into cooler times/places, we test whether heat lowers mixing in Chapters 7 and 8.

2.6 Synthesis

Framing this review, we treat cities as social reactors—places where payoffs turn on who meets whom, considering both planned and spontaneous encounters—and ask: as remote work and climate stress shift routines, can we still engineer durable connections and serendipitous interactions? The availability of large-scale GPS mobility data now makes these questions tractable across cities and over time, and the stakes are clear: misdiagnosis risks entrenching segregation, while sound measurement creates levers for inclusive growth. We organise what follows around three strands that motivate the thesis: (i) how experienced segregation and mixing arise through daily mobility; and how contemporary disruptions—(ii) hybrid work, the "introvert economy", and (iii) climate extremes—reconfigure incentives and change routines. To move from patterns to mechanisms, we distinguish mixing from bridging—deviations from a distance-and opportunity-based counterfactual—so we can see where diversity is produced rather than merely observed. This sets up the empirical programme: measure intergroup interaction from digital traces, map structure across scales, and test a simple rule—centres mix; peripheries sort—before tracing the consequences of decentralised work (the "doughnut effect") and rising heat. Our findings show that distant trips contribute a disproportionate share of intergroup exposure; weakening the centre reduces mixing; and hot days reduce, reshape and retime activity—often into the evening—altering the city's social reactor itself.

Segregation and Mixing

3 Rings and pockets of experienced segregation

Abstract As discussed, cities generate wealth from interactions, but citizens often experience segregation in their daily urban movements. Using GPS location data in this chapter, we identify patterns of this experienced segregation across US cities, differentiating between neighborhoods that are sources and sinks—exporters and importers—of diversity. By clustering areas with similar mobility signatures, capturing both the diversity of visitors and the exposure of neighborhoods to diversity, we uncover a generic mesoscopic structure: rings of isolation around cities and internal pockets of segregation. Using a decision tree, we identify the key predictors of isolation and segregation: race, wealth and geographic centrality. We show that these patterns are persistent across time and prevalent across all US cities, with a trend toward larger rings and stronger pockets after the pandemic. These findings offer insights into the dynamics that contribute to inequality between neighborhoods, so that targeted interventions promoting economic opportunity can be developed.

Introduction

The free flow of people, goods and ideas drives urban agglomeration. Cities enable "sharing, learning and matching" to generate efficient and innovative economies [171]: shared infrastructure, shared ideas, and pooled workers drive the dominance of urban systems in the economy today. By extension, cities that fail to integrate communities or struggle to connect residents may sacrifice growth [222] or exacerbate poverty [25, 144]. A large body of literature documents housing segregation by race and class, and recent developments, including mobile phone location data, enable us to investigate segregation in daily life as people move around the city [38, 321], revealing bias in who interacts with whom. Since urban innovation and prosperity hinge on diverse interactions [416, 68], understanding the nature, extent, and limitations of how groups interact is essential to understanding and building inclusive economies.

Cities in the United States exhibit "hypersegregation", the separation and concentration of nonwhite populations in contiguous zones, typically near the urban core [301]—often a result of discriminatory institutions [424]. While implicit in hypersegregation is isolation across multiple dimensions, GPS mobility data reveals that "experienced segregation"—wherein people are more likely to share spaces with others of the same race [38] and class [321]—is a key component, affecting crime rates and economic opportunities [271]. Yet recent work emphasizes that nearby venues can have distinct visitor profiles [321, 150], suggesting that sorting occurs beyond residential confines. Studies also indicate that amenity location can mediate interactions between groups [335, 336]. Mixing across strata often involves lower-class individuals visiting higher-class areas [231]—a pattern consistent with hypersegregation and "compelled mobility" [88], which involves disinvestment and consequent travel for basic services.

This matters because generating connections between groups and across communities may create wealth and reduce inequality. Early work on segregation implicates the lack of "exposure" between different groups in creating economic inequalities [302]. While work on experienced segregation measures exposure, not necessarily interaction, mobility links to social ties and their social and economic consequences. Activity spaces—the bounds of our routines in space—influence social connections [261, 81], with particular consequences for "weak ties" [211] that generate social and economic mobility. Employment opportunities within an individual's activity space predict job attainment [97]. Housing segregation correlates with low social capital and social mobility [118, 119], but recent work finds that areas with venues that encourage mixing also tend to have greater social capital and social mobility [299]. Neighborhood connectivity, including who visits and where residents go, predicts homicide rates—with the most disadvantaged communities those defined by both residential and experiential segregation [271].

Although the literature on experienced segregation is growing [277], recent research tends to focus on human behaviors [321, 522, 81, 231, 299, 132] or urban comparisons [38, 490, 3, 335]. These microscopic and macroscopic views often neglect the mesoscopic clusters of neighborhoods within and around cities that drive mixing or isolation, and the forces that generate these zones. Yet the mesoscopic level is where phenomena like hypersegregation [301] operate—creating not just isolated neighborhoods but interconnected systems of disadvantage. Limited connections to the broader economy create "truly disadvantaged" communities [512], but new connections ameliorate some of these problems [116, 115, 123]. The location and concentration of experienced segregation reveals not just who is isolated but also who is surrounded by isolation, hindering new connections.

Here we seek to identify the broader zones of mixing and isolation that define cities. How residents interact with the broader city reveals inequalities between groups, as certain groups or places may make uneven contributions to aggregate urban mixing. This work investigates the mesoscale of mixing in US cities with a spatially and temporally expansive sample of cell phone GPS records covering the continental United States over 4 years. We focus on visits to amenities as they represent both primary attractions where populations purposefully interact [335, 299] and structured environments where social connection occurs [456]. We present two key measures: amenity *segregation*, representing the degree to which visitors to an amenity are diverse, and neighborhood *isolation*, indicating the degree to which travelers from a neighborhood experience diversity in their day-to-day activities. Combined, these related measures reveal distinct patterns, including neighborhoods—and clusters of neighborhoods—that export diversity to the broader city without importing it. Our approach allows us to empirically identify clustering as a dimension of experienced segregation, revealing spatial structures that would be obscure at the scale of either the city or neighborhood.

Cities operate as social reactors that turn proximity into interaction, both planned and unplanned, and productivity; the rings and pockets we uncover show the nuances and variations within that social reactor, concentrating the meetings and exposures that sustain urban learning, sharing, and matching [68, 171].

Results

Data and methods

To assess experienced segregation across multiple dimensions, we construct measures for diversity and exposure in terms of socio-economic class. We process data from a lo-

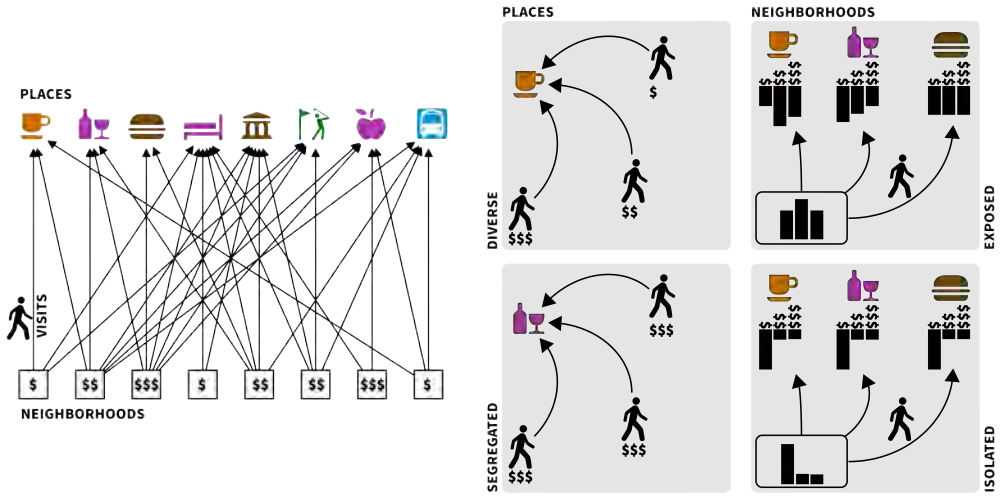

Figure 3.1. **Illustration of how we measure segregation and isolation**. Residents from neighborhoods with different median incomes visit amenities. Amenities are segregated if they attract visitors from a similar income bracket; neighborhoods are isolated if their residents visit segregated amenities during their daily activities throughout the city.

cation data provider [426] to create origin-destination flows from neighborhoods (census block groups) to points of interest (like restaurants, shops, museums, and hospitals) over a period from January 2019 through December 2022. We combine these data with Census estimates of median income [498] to infer the socio-economic strata of visitors from each area by attributing to each visitor the median income of the block group where they live. We compute the "diversity" of visitors to amenities and the "exposure" of neighborhood residents to diversity. Borrowing from prior work [321], our measures represent distance from a counterfactual scenario where interactions draw equally from each socio-economic stratum in a given city. We examine amenity segregation (low level of diversity of visitors to an amenity) and neighborhood isolation (low level of exposure of residents to diversity in the amenities they visit), considering an amenity highly segregated if it attracts visitors from just one income bracket, and residents highly isolated if they only visit such amenities. Both measures range from [0, 1], where 0 is perfect diversity or exposure and 1 is perfect segregation or isolation. Detailed in the Methods section, we illustrate our process in Fig. 3.1.

Segregation and isolation

We can think of a neighborhood as *isolated (I)* and an *amenity*, or point of interest (POI), as *segregated (S)*. These are illustrated in Fig. 3.2 for a sample of 10 large cities. While previous work [321] emphasizes variability between nearby amenities, showing that a pair of adjacent amenities may have opposite diversity signatures, Fig. 3.2A shows that these form distinct clusters of both segregation and diversity. Our sample includes cities with different spatial structures, from sprawling "sunbelt" cities like Houston and Dallas to dense "rustbelt" cities like Chicago and Philadelphia. The latter tend to have pockets of segregation, like North and West Philadelphia, the south side of Chicago, and the South Bronx in New York. In contrast, Houston and Dallas are more integrated. Cutting against this distinction, Los Angeles has a large pocket of segregation in South Los Angeles—a historically disadvantaged neighborhood. Prior work has shown that hypersegregated populations have historically been located near downtowns, confined to contiguous areas [301, 302], much as we observe here.

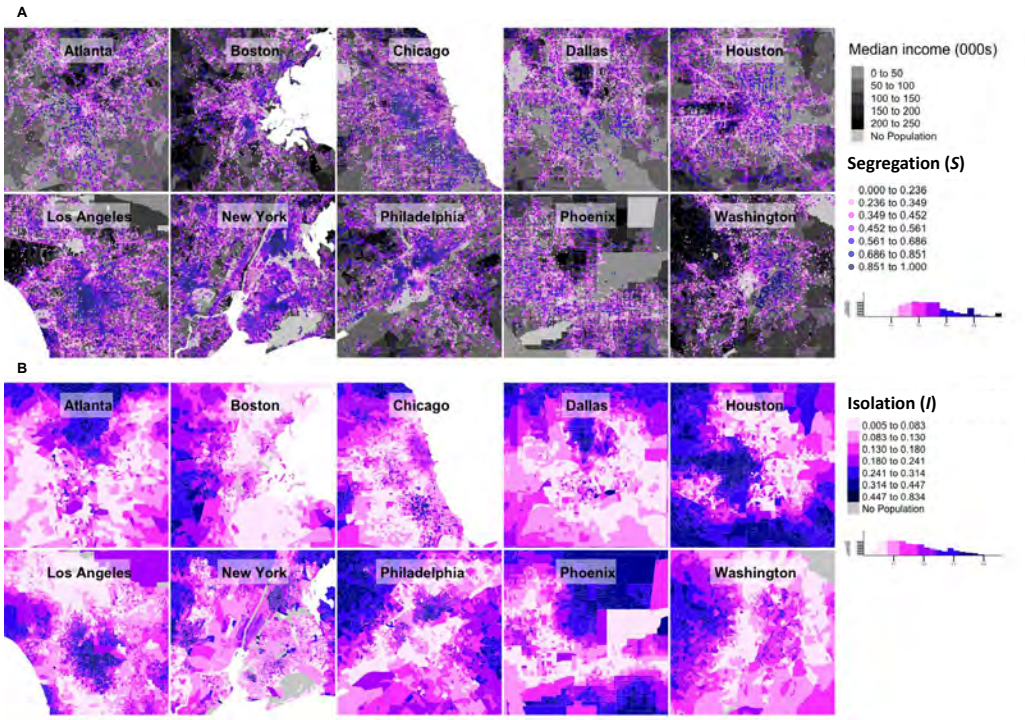

Figure 3.2. **Illustration of amenity segregation (*S*) and neighborhood isolation (*I*).** Each city map is centered on downtown. **A** Amenity segregation, where each point represents an amenity, shows that downtown businesses tend to see a diverse collection of visitors (and thus have low amenity segregation) but that businesses in surrounding neighborhoods often do not (and have high amenity segregation). Many of the wealthiest parts of the city, shown in black, also have fewer points of interest, which limits visitation and thus the diversity. **B** Neighborhood isolation is strong in those same wealthy areas with fewer amenities and also in areas with segregated amenities.

One channel by which segregation is determined is in the mix and number of amenities in a cluster: fewer reasons to visit will limit visitors. Many wealthy areas—shown in black—simply have fewer amenities, and less variety of them. Showing this correlation in Appendix Table A.3, we see examples of this in Dallas, Phoenix, Washington, and Los Angeles—corresponding with Bel-Air and other wealthy enclaves in Fig. 3.2A.

To understand the relationship between segregation and isolation, we bin each by quantile and construct a 3x3 matrix that classifies each neighborhood by its performance on both measures. A neighborhood that is completely integrated, with diverse amenities and with residents visiting diverse amenities, will be in the first quantile along both dimensions, segregation and isolation. Vice versa, homogeneous amenities and residents visiting homogeneous amenities will be in the third quantile for both. Residents may also find exposure to diversity away from home, and neighborhoods with them will be in a low quantile for isolation and a high quantile for segregation. Further, because we capture segregation at the level of the amenity, we capture sorting into amenities with a given demographic profile even as the neighborhood around them receives a different selection of visitors.

In Fig. 3.3A we find that many cities are surrounded by rings with high segregation and high isolation. In the Boston-Washington megalopolis, many of these rings merge into large intercity zones. Looking within cities, we see that there are areas where visitors are diverse but exposure is low (aqua colored areas in Fig. 3.3B): residents sort into amenities that allow them to avoid ambient diversity. These neighborhoods are *heterophobic*, and they are typically central and affluent. The opposite also occurs, often in the same cities (pink colored areas in Fig. 3.3B): *heterophilic* neighborhoods in Chicago and

New York have residents that experience diversity despite living in a neighborhood that sees little of it. We again see that historically deprived neighborhoods, like the South Bronx or South Los Angeles, are often both segregated and isolated. To understand these neighborhoods, we plot distributions of selected variables by class in Fig. 3.3C (expanded to include amenity distributions in Fig. A.3): neighborhoods which export diversity without importing it tend to exist at a certain distance from downtown, with characteristic socio-economic attributes—high nonwhite populations and low incomes. They also tend to have fewer amenities, while areas with residents who appear to avoid diversity have more—possibly enabling sorting.

Many of the urban *pockets* we see have been victims of federal, state or local policies to produce segregation—or at the very least a policy failure to remedy segregation [424]. In Appendix Fig. A.4, we show that areas that were given poor grades—usually due to prejudice against nonwhite residents—by the Home Owners Loan Corporation in the 1930s still exhibit higher levels of segregation today, and areas that were not graded at all—because cities had not sprawled out into the suburbs until later in the century—exhibit higher levels of isolation. These ungraded areas are likely to be the very places that excluded nonwhite families. This suggests that some of what we see in the data is the result of policy. We document structural developments that may have influenced certain pockets in Appendix A.5.

We include sensitivity analyses and robustness checks in Appendix A.6 and A.9. We also show that neither home nor work location are likely to be the cause of these patterns, suggesting at least some degree of sorting on preference or on price, in Appendix Fig. A.11.

Zones of segregation and isolation

To understand how these measures of interaction manifest in the structure of cities, we use tests for spatial autocorrelation to identify contiguous zones of amenity segregation or neighborhood isolation. These tests show that at both local and regional scales there are large areas where high values cluster. Segregation and isolation exist at characteristic scales, with the latter clustering in suburban rings, Fig. 3.4A, and the former clustering in certain urban pockets, 3.4C. We take the top 100 cities and subtract the central business district (CBD) coordinates from each so they stack on top of each other (see Methods). In Fig. 3.4B we show the composite pattern by counting the number of isolated clusters from all cities at each location relative to their centers. This reveals the relationship between centrality and isolation, showing that these rings form at similar relative distances across different cities regardless of geography. Compared to clusters of isolation, which tend to be suburban, segregation tends to be urban and the clusters tend to exist at a smaller, more fragmented scale. Fig. 3.4C shows that large cities—New York, Los Angeles, or Chicago—also tend to have larger concentrations of segregation, compared to small cities like Atlanta and Boston.

Isolated populations in these zones scale predictably with city size, shown in Fig. 3.4D. For isolated clusters, the relationship is superlinear ($\beta = 1.27$); larger cities present larger concentrations of isolation. The population in segregated clusters scales superlinearly ($\beta = 1.37$) with the population of the city. As cities become larger, separated areas with limited visitation become larger at a faster rate. Large areas visible in big cities (New York, Chicago) and less so in smaller cities (Atlanta, Boston) in Fig. 3.4C.

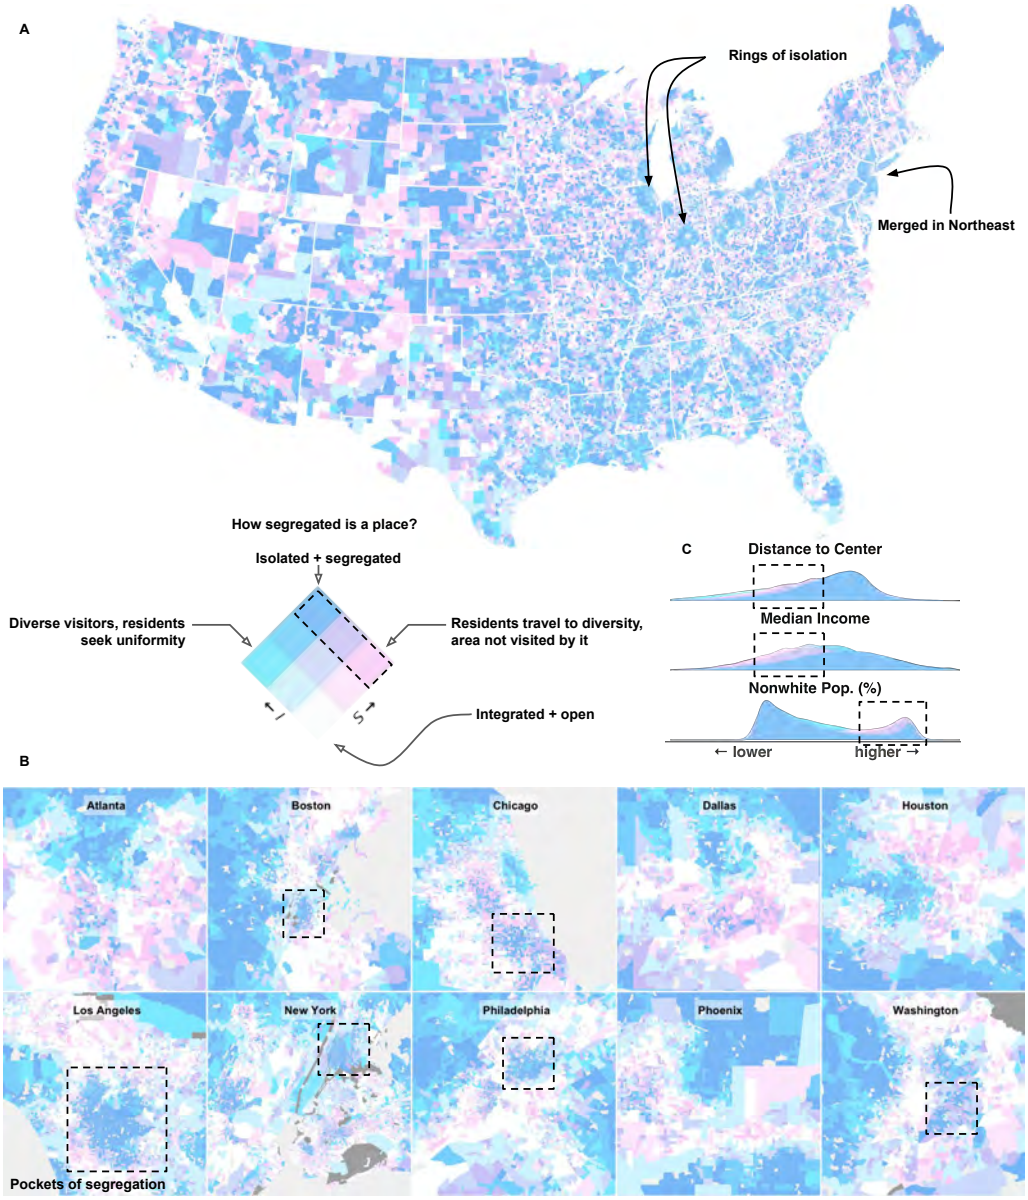

Figure 3.3. **The relationship between segregation and isolation. A** Mapping diversity and exposure nationally we see that integrated urban areas are often surrounded by rings of isolated suburban areas. **B** Locally, we see urban pockets of segregation with a range of low to high isolation, often near or surrounded by integrated urban areas. **C** Distributions of selected variables by class show that segregated areas are often close to the center, poorer than average, and nonwhite.

Trends and drivers

We use a decision tree, shown in Fig. 3.5A, to understand what defines each class of our 3x3 scheme. We append attributes to each neighborhood using data from the Census to serve as predictors. We also measure the distance to the CBD, identified earlier. These become predictors in our model, which provides us with an intuitive, step-by-step way to see which variables predict which class in order of importance. Points of decision include richer or poorer, densely or sparsely populated, predominantly white or nonwhite, centrally or peripherally located. For example, we can ask what is the most common class of a neighborhood that is near the center, predominantly nonwhite, relatively dense and relatively poor.

Our complete tree, shown in Appendix Fig. A.12, achieves an accuracy of 30%, a

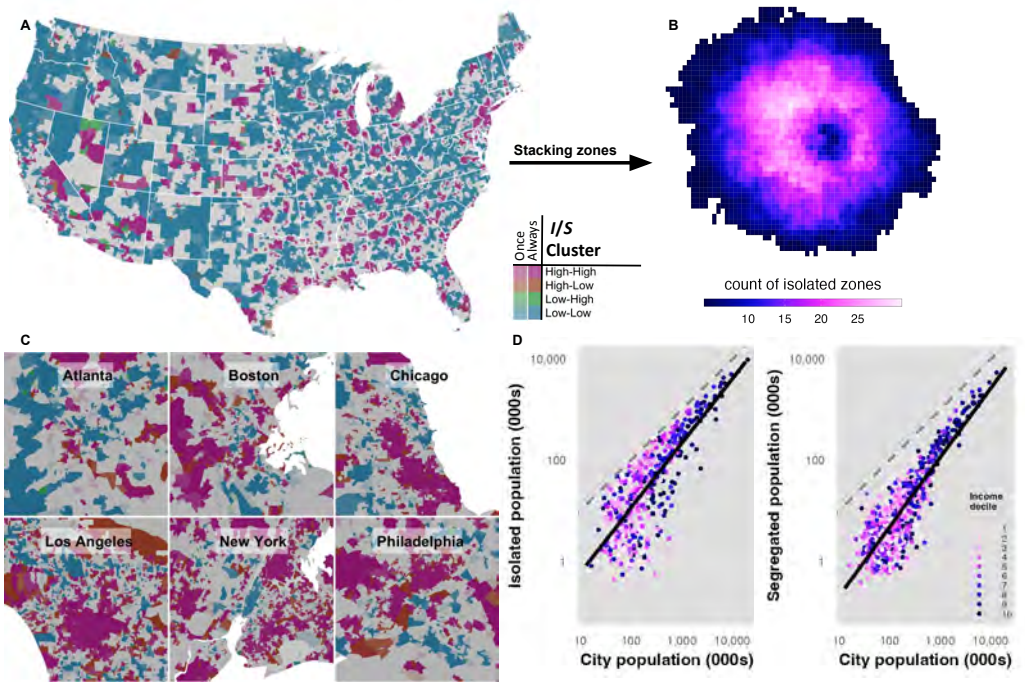

Figure 3.4. **Defining rings of isolation and pockets of segregation. A** Isolation autocorrelation manifests at the national scale, delineating rings around cities. Centering and layering the cities, in **B** we count the number of times isolated zones occur in the same relative area: there is a clear prevalence of these zones in a ring surrounding each urban core. **C** Segregation autocorrelation manifests locally, with pockets appearing in large cities—less so in smaller cities like Atlanta or Boston. We also see tight scaling relationships between city size and isolation/segregated population in **D**, with a linear relationship between city population and the population in isolated zones, along with a superlinear relationship for city population and its segregated zones.

threefold improvement across what would be expected if assigning classifications from the 3x3 matrix to neighborhoods at random. Appendix Table A.5 shows the importance of each feature in the tree, using a permutation technique that shuffles each column and computes the resulting reduction in predictive ability. Income, along with measures of urban structure like density and distance from center, are the most important predictors. This ranking suggests that, along with income, urban form and function, including the distribution of amenities, may be important determinants of segregation and isolation. Proximate branches a pruned tree in Fig. 3.5A indicate that both the wealthiest neighborhoods and dense, nonwhite neighborhoods are the most segregated. The most likely populations to have high segregation and low isolation are nonwhite. In less affluent and less white communities, the presence of amenities is associated with a greater diversity of visitors.

The data suggest employment is not important. Insofar as a job represents a connection to the broader urban economy, we might expect the size of the working population to predict isolation above and beyond income. Instead, the best model uses income and race, along with spatial structure, to estimate isolation and segregation. This fits with the weak correlation between commutes and mixing: where you work does not do much to help exposure if you sort into venues with an audience similar to you.

The decision tree allows us to explore the relationship between key variables conditioning on all other predictors [204], modeling the interactions in the data. Fig. 3.5B and C show the model's estimated values for segregation and isolation, respectively, given certain conditions; it asks, given these values for demographic, economic or geographic features, what do we expect the segregation and isolation to be?

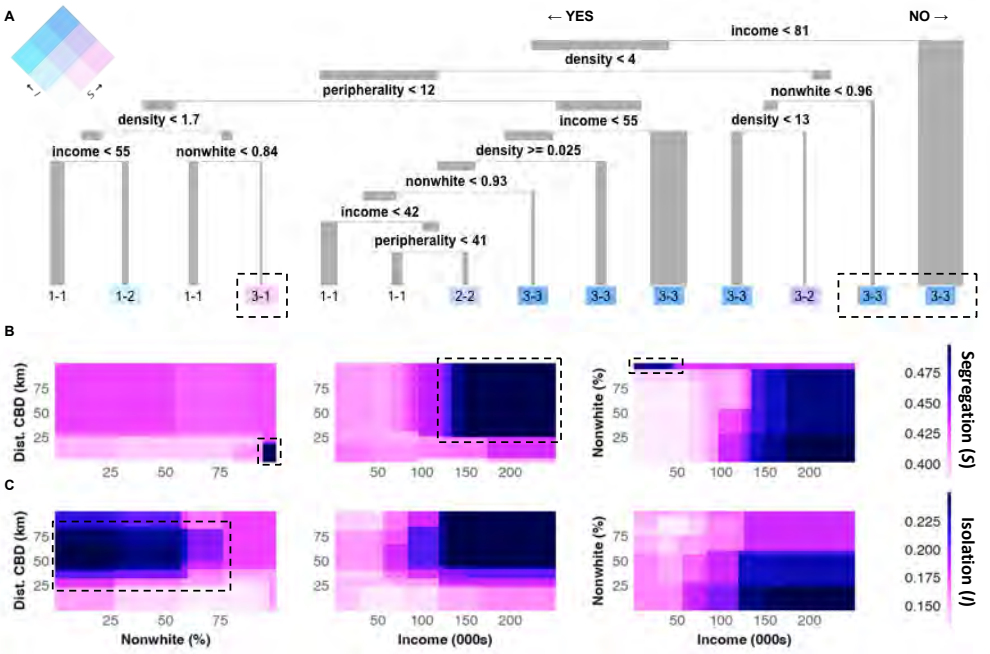

Figure 3.5. **Factors associated with segregation and isolation. A** Decision tree showing the defining characteristics of different classes, pruned for ease of viewing. Note that top segregated and isolated classes are the wealthiest cut, but also urban nonwhite—as indicated by density. The areas that have high segregation but low isolation tend to be urban, nonwhite and moderately dense. Partial dependence plots for **B** segregation and **C** isolation showing the joint relationship between key variables, with rings of isolation conditional on white/wealthy and pockets of segregation conditional on nonwhite/poor.

The patterns are clear: if a neighborhood is close to its downtown, its amenities are more integrated, unless it is majority nonwhite; if a neighborhood is far from its downtown, it is more isolated, unless it is majority nonwhite. The relationship between location and demographic composition creates quadrants that correspond to higher or lower segregation. The tree indicates that race is the most important driver of amenity segregation: the highest modeled estimates for amenity segregation come in neighborhoods with large nonwhite populations—regardless of how central they are. Wealthy neighborhoods also have more segregated amenities, but only if they are majority white. The rings around each city are also evident in these distributions, with segregated suburbs becoming comparably less segregated exurbs—which are also less expensive–at the farthest extents of the city.

Given the potential consequences for segregation and isolation, we link these measures of spatial mobility to social mobility—the likelihood that a child will earn more than her parents—in Appendix Figs. A.13 and A.14. In brief, higher segregation and isolation correspond with better outcomes, suggesting strong selection effects—except for areas with more than 75% nonwhite residents, where the effects go negative. This indicates that connections to the broader city are important for some groups but not others.

Temporal variations

Across time, the bivariate class comprised of the most segregated and isolated areas, flowing across the top of Fig. 3.6A, became larger during the pandemic but has since returned to its 2019 fraction of the total. The class comprised of the least segregated amenities, flowing across the bottom, has not recovered—an indication that downtowns are less places of mixing than before. Despite this churn, transition probabilities are

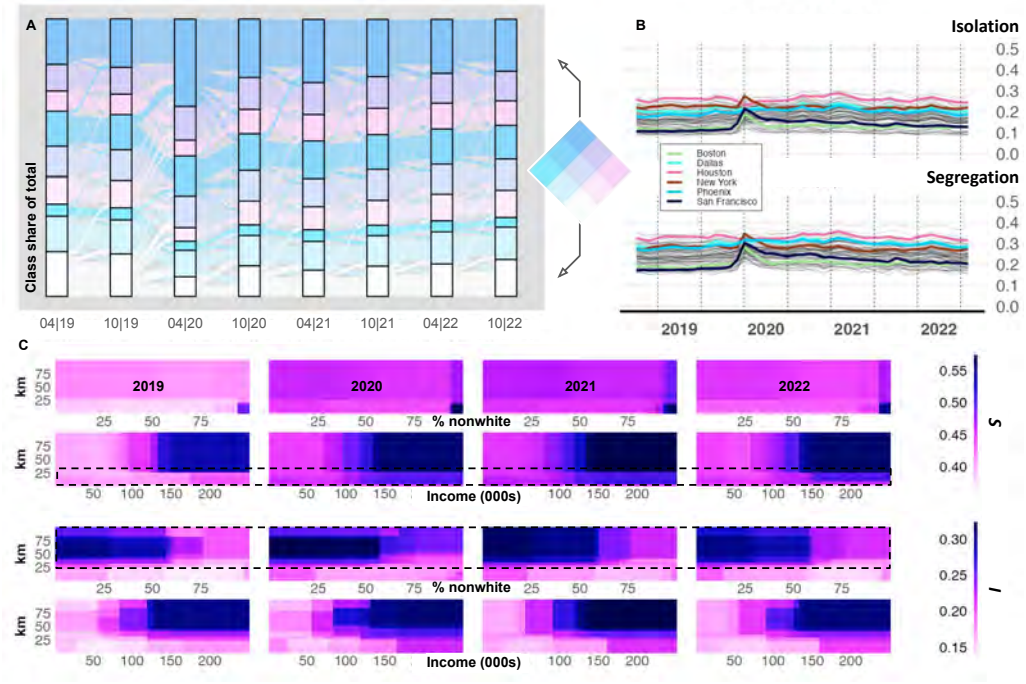

Figure 3.6. **Changes in factors over time. A** Trends for classes of segregation showing that segregation peaked in April of 2020, but shows little change today. We compute city averages for I and S in **B** and see that levels changed in many cities during the pandemic but have returned to normal in most—notably not San Francisco. **C** Changes to dependencies over years, following the same logic as above with distance to the CBD on the vertical axis and variables of interest on the horizontal axes (nonwhite population and income): relative to the 2019, lower values for nonwhite population and median income generate high expected segregation near the center; more distance bands generate high expected isolated even at higher values for nonwhite population and lower values for median income.

low at the extremes, which we show in Appendix Fig. A.15: most and least integrated areas typically remain so. Confirming previous work [522], experienced segregation peaked during the pandemic; we extend this work through 2022, showing that levels of segregation and isolation have returned to 2019 in many cities. Computing the average values for I and S per city in Fig. 3.6B shows that, aside from San Francisco and Boston, the distribution has largely returned to its index value.

In order to understand how this change manifests in the structure of cities, we track how the partial dependencies have shifted over time in Fig. 3.6C. Most predictions, either for segregation or isolation, shift to higher values throughout the study period. A notable change occurs in the band representing wealthy neighborhoods close to the urban core, which become more segregated in 2021 and 2022 despite proximity to downtown amenities that facilitated mixing in 2019. Our 2019 model expects urban areas with 90% nonwhite populations showed high amenity segregation prepandemic. In the 2022 model, this dropped to 80%. Urban areas 5-10km from city centers with $200,000 median incomes saw segregation indices rise from $S = 0.4$ in 2019 to $S = 0.5$ in 2022—a 20% increase. The rings that we see around cities above also extend out: predictions for isolation become higher beyond 75km, where the highest predictions ended in our 2019 baseline. In these exurban areas or satellite cities greater than 75km from the city center, the average estimated isolation rises from $I = 0.25$ to $I = 0.35$ where neighborhoods are majority white.

We also see little change in the rank importance of features over the period analyzed here (see Appendix Fig. A.17). Although it maintains its rank importance, the level of importance of distance to the CBD does drop in 2021 and 2022, suggesting that we may

be seeing changes to the spatial structure of segregation and isolation.

Discussion

While prior work highlights subtle variations between adjacent venues [321, 150], here we show that cities often have a structure that minimizes mixing in some areas while maximizing it in others. Our analysis reveals that cities are comprised of spatially contiguous zones of isolation and segregation, where residents primarily interact with socioeconomically similar others. External rings can often stretch for great distances, encompassing large populations. Internal pockets, those exhibiting high segregation and either high or moderate isolation, resonate with concepts of the "ghetto" [168], a spatial arrangement that allows economic participation while maintaining social separation. The mesoscopic view then advances our understanding of experienced segregation by documenting the twin experienced concentration of affluence and poverty, opportunity and disadvantage, building on earlier work on socio-economic extremes in cities [300].

Studies documenting experienced segregation often fail to distinguish areas where diverse encounters happen from areas where residents must travel to encounter diversity [38, 321, 335]. Recent work highlights the"compelled mobility" of nonwhite residents who often travel outside segregated home neighborhoods to access resources [88]. Our findings align with this, showing some segregated areas (like South Chicago) whose residents experience diversity elsewhere, yet also identifying isolated pockets (like South Los Angeles) where both residential and experiential segregation coincide indicating disadvantage [271]. The other side of this are isolated areas (like most suburbs) where residents avoid diversity. While exposure offers social benefits [116, 299, 123, 118], travel for amenities may disadvantage businesses in isolated and segregated communities.

A key limitation of our study is its focus on the United States, but our findings have broad implications for policy in American cities. Research shows that "disadvantaged" communities [512] often have limited connections to the surrounding labor market [113]. Long ties in particular allow individuals to reach beyond their immediate communities, and these ties are grounded in activity spaces [243, 457, 261]. We identify disconnected neighborhoods that are in turn clustered with disconnected neighborhoods; individuals living in these areas are less likely to interact with the surrounding city *and* less likely to interact with someone who is connected to that broader economy, because the entire community in which they exist is disconnected. Constructing long ties will require connections between rich and poor communities—both the rings and the pockets in our study. Interventions should leverage zoning and land use to develop amenity clusters in accessible locations between zones [335], while investing in downtowns to restore them as places of mixing.

Although our topic intersects with racial segregation in the United States, our mobility data do not include information on race/ethnicity. To avoid ecological misclassification, we do not impute race/ethnicity from neighborhood composition and restrict our primary segregation metrics to income. Future work that links mobility to validated individual demographics is needed to assess racial/ethnic patterns and inform policy.

COVID-19 and remote work are rewiring urban interaction networks [380, 53]. Research documents pandemic changes in mobility generally [433] and experienced segregation specifically [522], finding that COVID-19 resulted in considerable changes to daily routine and to socio-economic mixing. The office [523] and the commute [316] are key determinants of "third places" [342]—shops, restaurants and cafes—we visit throughout the day. Further, a postpandemic "introvert economy" [447], with fewer nights out, has developed. Our results indicate that urban interaction patterns continue to evolve.

Changing work and leisure dynamics may have important consequences for social mobility, social capital and the relative economic advantage and disadvantage of neighourhoods. These dynamics may reinforce both concentrated advantage ("rings") and concentrated disadvantage ("pockets"), making connections between these extremes even harder to establish—potentially exacerbating the "age of extremes", marked by concentrated poverty and affluence [300]. Understanding these shifts is critical for creating opportunity, as interventions need to adapt to new mobility and interaction realities.

In the following chapter, we extend this analysis by looking and the points of interest that modulate interactions between groups. We introduce the concept of "bridging": locations that are situated between communities with different racial or socio-economic compositions, encouraging members of both to meet in the middle. The intuition for why meeting places matter has its roots in the intervening opportunities models that we discussed earlier, but the stacked "bullseye" of isolation and segregation—a centre for mixing, adjacent zones of segregation followed by rings of isolation—prefigures why bridging amenities are not the most important mechanism for driving mixing: when central locations dominate travel budgets, intervening opportunities are leapfrogged. Put simply, this chapter maps where the social reactor runs cool; our next chapter explains where and why it runs hot even when the geometry of bridging predicts otherwise.

Methods

With the goal of assessing experienced segregation along multiple dimensions, we construct measures for diversity and exposure for both socio-economic class and race. To achieve this, we use data from SafeGraph [426], a location services provider, to construct origin-destination flows from neighborhoods (Census block groups) to points of interest. SafeGraph gathers GPS locations from mobile phones by aggregating data from applications who have obtained user consent to passively monitor location, creating a sample of users comprising $\sim 10\%$ of the population. Their process assigns visits to points of interest by clustering GPS pings and joining these clusters to adjacent building polygons, using relative distances and time-of-day to manage conflicts [426]. The data have been used in a variety of academic contexts [109, 106], and we validate the data in Appendix Section A.1 to show that the number of devices in a block group shows a strong correlation with population and no detectable correlation with income or race, which would indicate systematic bias. The result is a rectangular matrix consisting of 220,000 origin home block groups (a Census aggregation with a population ~ 1000) and 7 million destination points of interest like restaurants, museums, cafes, car dealers, or grocery stores. To each origin we append Census estimates from the American Community Survey [498] of median income and nonwhite population in order to infer socio-economic strata and demography of the visitors from that area. With these we can compute the "diversity" of (visitors to) points of interest and the "exposure" of (residents in) neighborhoods to diversity.

Measuring intergroup interaction. We look at amenity segregation and neighborhood isolation. The first asks, how diverse are visitors to this place? The second asks, how much diversity are residents from this neighborhood exposed to in daily routine? Segregation captures the patrons at a given amenity, or point of interest (POI), which we can then use to compute isolation at the level of the neighborhood by taking the weighted average of segregation at points of interest visited by the residents of a given neighborhood. The measures we construct represent deviations from a counterfactual scenario where visitors to an amenity draw equally from each socio-economic stratum in a given

city. If a restaurant attracts visitors from just one income bracket, calibrated within that metropolitan area, we consider that amenity to be highly segregated; if residents from a neighborhood only go to restaurants like that one, we consider them to be highly isolated. Diversity and exposure are the inverse of segregation and isolation. We illustrate this process in Fig. 3.1.

Conceptually, segregation is inversely related to the diversity of visitors at an amenity (low segregation S means high visitor diversity), while isolation is inversely related to residents' exposure to diversity during their activities (low isolation I means high exposure to diverse amenities). While mathematically related through the calculation of isolation from segregation, they capture distinct aspects of urban mixing: the nature of individual places versus the aggregated experience of neighborhood residents.

Following earlier work [321], we consider the segregation S of an amenity α to be a distance from an ideal scenario where people from all socio-economic classes visit in equal proportions. This is defined as follows

$$S_\alpha = \frac{5}{8} \sum_q \left| \nu_{q\alpha} - \frac{1}{5} \right|, \tag{3.1}$$

where q represents an income quintile and ν represents the portion of visitors from that quintile. We scale that by $\frac{5}{8}$ so that each value spans 0 to 1, with 0 being perfect integration (equal proportions from all classes) and 1 being perfect segregation (visitors from a single class). Each quintile is calibrated to the metropolitan area, rather than the nation as a whole. As a robustness check, we also compute quintiles based on the nonwhite population, differentiating between neighborhoods based on the proportion of nonwhite residents.

An inherent limitation of this study is that we are assigning these classes according to area, rather than individual attributes. This leaves the possibility for misclassification, which we address through sensitivity analyses in Appendix A.6, where we consider how our results would change if devices were drawn from different parts of the income distribution of each block group, but we cannot be certain about the quality of our imputation. Because income is ordinal and equally partitioned while race/ethnicity is nominal with unequal base rates, our mixing measure fits income better; for race/ethnicity, ecological assignment and category collapse make the measure both noisier and more fragile in interpretation, so we treat those results as complementary rather than primary.

Neighborhood isolation I is obtained by aggregating S as follows. We compute the average diversity S for all amenities α visited by residents of a given neighborhood γ, weighted by the number of trips T between neighborhood γ and amenity α, leading to the following definition

$$I_\gamma = \frac{\sum_\alpha (S_\alpha \times T_{\gamma\alpha})}{\sum_\alpha T_{\gamma\alpha}}. \tag{3.2}$$

We then use Local Indicators of Spatial Autocorrelation (LISA) [29] to show the existence of larger zones of segregation and isolation, which allows us to understand how cities are structured. Spatial autocorrelation refers to the degree to which similar values cluster together in space, and this test lets us identify the clusters of contiguous zones where residents have limited exposure to diversity. LISA measures the correlation between an areal unit and its neighbors along some dimension; if values of segregation or isolation co-occur in space, we will see clustering. We use permutations, shuffling the values of our variables of interest and recomputing autocorrelation, to test the significance of each local I value, keeping only the values for which $p < 0.05$. The results can be used to not only to detect clusters of similar values, which we define as "high-high" when the

index cell and its neighbors are in the top third of segregation or isolation; we use “low-low" for the opposite, as well as the remaining combinations of high and low mixes.

To understand how regular and predictable these clusters are across cities, we then perform a transformation to align cities. We identify the “downtown” by finding the largest cluster of amenities within 250m of each other using DBSCAN [219], which groups points based on spatial density. We calculate the centroid of this downtown cluster by averaging all amenity coordinates, weighted by visit frequency. With this urban centroid identified, we subtract this X,Y value from all geometries of all block groups in a city. This has the effect of moving all cities to a common centroid around 0,0—stacking them on top of each other. From there, we simply count the number of high-high clusters of isolation in each part of the city using a gridded mesh to standardize the block groups. The resulting heatmap reveals the spatial distribution of isolated zones relative to city centers across our entire sample.

Dimensionality reduction. Our data span four years, from 2019 through 2022, which gives us the ability to assess the stability of these measures over time. To understand this, we both follow the raw trends and track relationships for a series of variables that predict diversity and exposure. We first build a time series of segregation and isolation for each neighborhood in our data, 48 observations (months) across 220,0000 neighborhoods (all Census block groups). Monitoring every time series presents a challenge, so we divide our data into 9 classes—combining 3 quantiles of segregation with 3 quantiles of isolation—at 2019 January t_0, and then assign classes according to the original breaks for each month t_i thereafter, tracking the size of each class throughout our sample. We also compute average segregation and isolation for each metropolitan area and track those levels as well. Because our data represent a space-time cube, where each location in space has 48 values, we explore other forms of dimensionality reduction in Appendix A.14.

To understand the factors associated with these trends, we construct decision trees to predict segregation and isolation at each interval. We choose this modeling strategy because the patterns in our data involve symmetries and nonlinearities, with high and low extremes of a given variable giving similar predictions, and because variables appear to interact. A decision tree manages these challenges in a manner that preserves interpretability, as we can explore the tree to see how variables are stacked to generate predictions [86]. The key tools that we use to interpret our data are the tree itself but also partial dependence plots, which show the best prediction given the values of two variables controlling for all others [204].

Our decision tree uses the following predictors, all at the level of the block group, from the Census Bureau (see Table 3.1): median income, nonwhite population share, college educated share of adults, household size (Census Bureau), rental vacancy rate, share of rent burdened households, share under 16 years old, share unemployed, and population density. We then use data from SafeGraph to construct the following metrics, also at the level of the block group: amenity count, which is the number of POIs in the block group, amenity entropy (H), which is the diversity of those POIs according to the 6-digit NAICS industry classification (E.g. Restaurants and Other Eating Places). We compute it as Shannon entropy [451] $H = -\Sigma p(x) log p(x)$ for all x where $p(x)$ is the proportion of firms in each industry classification. We also use DBSCAN [219] to identify each city's centroid in the following way: we sort clusters of POIs within 50 meters of each other by the number of POIs in them, and assume that the largest cluster is the central business district; we then take the visit-weighted average latitude and longitude over those points of interest and assume that point is the city center. We then measure the distance of each block group in each city to that point and call that

distance to the CBD, which we use as a predictor.

Variable	**Definition**
Dist. to CBD	Distance in kilometers from the block group to the central business district; CBD identified as the largest cluster of POIs within 50 meters of each other using DBSCAN algorithm
Amenity (H)	Shannon entropy of POI types, calculated as $H = -\Sigma p(x) \log p(x)$ where $p(x)$ is the proportion of POIs in each 6-digit NAICS industry classification
Density	Population density (population per square kilometer)
Amenity (#)	Count of POIs
Median income	Median household income in US dollars
Nonwhite (%)	Percentage of population not identifying as non-Hispanic white
College (%)	Percentage of adults (age 25+) with at least a bachelor's degree
Household size	Average number of people per household
Vacancy rate	Percentage of rental units that are vacant
Rent burden	Percentage of households spending more than 30% of income on rent
Under 16 (%)	Percentage of population under 16 years of age
Unemployed (%)	Percentage of labor force that is unemployed

Table 3.1. **Definition of variables used in the decision tree model.** The model uses a combination of demographic, economic, and urban structure variables to predict neighborhood segregation and isolation. Data come from the US Census Bureau's American Community Survey and SafeGraph location data. All variables are measured at the Census block group level.

4 The role of bridging in mixing

Abstract Having discovered important patterns related to how different parts of cities integrate with the broader whole, in this chapter we show that “exposure segregation"—the degree to which individuals of one group are exposed to individuals of another in day-to-day mobility—is dependent on the structure of cities, and the importance of downtowns in particular. Recent work uses aggregated data to claim that the location of amenities can inhibit or facilitate interactions between groups: if a city is residentially segregated, as many American cities are, then amenities *between* segregated communities should encourage them to mix. We show that the relationship between “bridging" amenities and socio-economic mixing breaks down when we examine the amenities themselves, rather than the urban aggregates. For example, restaurants with locations that suggest low expected mixing do not, much of the time, have low mixing: there is only a weak correlation between bridging and mixing at the level of the restaurant, despite a strong correlation at the level of the supermarket. This is because downtowns—and the bundle of amenities that define them–tend not to be situated in bridge areas but play an important role in drawing diverse groups together.

Introduction

New research finds that the distribution of amenities in a city relates to socio-economic mixing: cities with “hubs" situated between communities of different socio-economic strata tend to have more exposure across those strata; cities with hubs located comfortably within homogeneous zones have less [335]. Using data from the same provider that we used in the previous chapter, the following analysis shows that although hubs are associated with mixing across metropolitan areas, when we sharpen the spatial resolution of the analysis, the association attenuates considerably, raising concerns about causality. Further, this relationship ignores the importance of hierarchy in urban structure first articulated in Central Place Theory, an idea which began in economic geography and has found support in recent studies of mobility. In brief, this theory posits a hierarchy of locations wherein some goods or services are provided frequently through dispersed vendors and others—central goods—are accessed infrequently through centralized agglomerations of amenities [122]. We can improve our understanding of mixing—and with it our policy recommendations—by first incorporating and then validating ideas on central places and central goods. In the following chapter, we show that the places doing the heaviest mixing work are not necessarily peripheral “bridges”, but central agglomerations—downtown clusters that mix despite low bridging potential.

Without explicit consideration for the role of central places and central goods in driving mobility, a growing body of literature indicates that the location and co-location of amenities in space is the product of related forces. Consistent with the idea that certain goods are traded centrally and others are served in a dispersed manner, with consequences for how often transactions take place, data on mobility show that visits to an area are a function of frequency and distance: some places are visited often by those nearby, and other places are visited sparingly by people from far afield [444]. Other work has re-

vealed distinct hierarchies in travel, with trips bounded by local or regional limits [17], and constraints on the number of places any given person can visit [18]. Together, much of the work on human mobility in the urban context suggests that individuals must allocate finite resources—time, money and even cognitive bandwidth [83, 47]—to satisfy needs. “Trip chaining"—whereby people visit multiple services on the same journey—is a common strategy for coping with these limitations [316]. Taken together, this research suggests that urban structure constrains human behavior, and vice versa—with amenities agglomerating to meet certain needs. Here we examine the implications of this interplay between structure and mobility on socio-economic mixing.

Results

We use a measure of mixing borrowed from earlier work [321] that captures the degree to which the distribution of visitors to a place deviates from an ideal wherein visitors from each quintile of the income distribution visit that place in equal proportions. These quintiles are calibrated using the metropolitan area, rather than the nation as a whole, to compensate for the concentration of wealth in some cities. The result is a value from 0 to 1 where 0 represents perfect mixing and 1 represents perfect segregation. With that as our proxy for mixing computed for ~ 5 million points of interest nationally, we construct our bridging index—which represents mixing under counterfactual mobility scenario—using a modified version of an established model [454] that assigns trips between origins and destinations considering the distance between them and the intervening opportunities; in this model, the probability that you visit a restaurant is a function of how far away that restaurant is and how many other restaurants are closer. We can think of this as an “expected” mixing to which we can compare observed data. We resample the observed data according to these probabilities for many business types. The above mixing measure, computed on these assigned trips, constitute our bridging index; this counterfactual mixing indicates the potential for amenities to attract diverse patrons.

Replicating earlier bridging-mixing findings [335], we find that our measures are well calibrated: bridging at the level of the city is correlated with mixing at the city as a whole ($\rho = 0.63$) and each of our measures for bridging and mixing are correlated with those of the earlier study ($\rho = 0.55$ and $\rho = 0.57$, respectively). Yet this relationship varies when we disaggregate to level of the point of interest: some amenities encourage mixing and others do not. Looking at restaurants, the Pearson correlation coefficient for mixing and bridging falls to 0.40 when we consider restaurants in Fig. 4.1**A**. At supermarkets, however, which are fewer in number and perhaps less separable according to taste and preference, the correlation coefficient between bridging and mixing is 0.70 (see Fig. 4.1**C**). Convenience stores are between these 2 extremes, with a correlation of 0.53 (see Fig. 4.1**B**). There is a notable surplus of restaurants in particular off the diagonal, where we would expect less mixing than occurs.

This presents a puzzle that calls into question the causal effect of bridging on mixing in American cities: cities with bridging amenities experience less exposure segregation than those without, but many amenities with low bridging also have high mixing. Our data focuses on venues, so it is plausible that much of the mixing in cities occurs on the street in front of restaurants and grocers as the patrons themselves sort into amenities along socio-economic characteristics. This sorting would align with work showing sorting among nearby venues across cities [321], and would indicate that exposure is superficial—shared paths rather than amenities. It is also plausible that unobserved variables produce both bridging and mixing, and that the two are not causally linked. We check our hypothesis that macro-mixing occurs alongside micro-sorting by aggregating venues

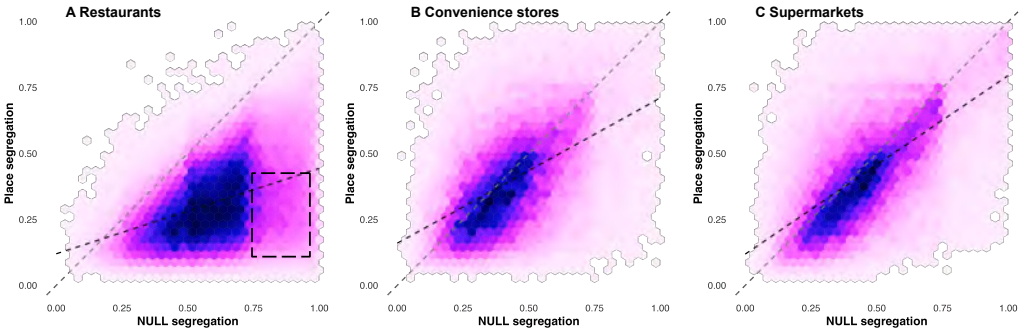

Figure 4.1. **Mixing and bridging.** Plots showing the correlation between bridging and mixing at the level of amenity for different venue classes, with weaker correlation for **A** restaurants and stronger correlation for **C** supermarkets, and **B** convenience stores in the middle. A large portion of restaurants are high expected segregation and low observed segregation.

into clusters and comparing cluster-level mixing to amenity-level mixing; we also check for systematic biases in the relationship between bridging and mixing. Fig. 4.2 presents these results.

When we subtract expected segregation from observed segregation, we see that the lowest values—segregated well below expectation, given location—are concentrated downtown (see: Fig. 4.2**A**). Further, the relationship between bridging and mixing breaks down when considering these business districts: these hubs typically have high predicted segregation—because they are situated far away from much of the broader suburban and exurban population—and low observed segregation. In Fig. 4.2**B**, we compute hub mixing across all clusters of amenities and then extract downtowns, which we define as the largest cluster in each city, to demonstrate. This also explains the weak relationship between mixing and bridging for restaurants: many restaurants are co-located in larger business clusters; these restaurants tend to be the low-bridging, high-mixing venues we see above. Considering the 2 large clusters in New York City and Chicago in particular, this relationship is partially driven by sorting within the cluster at the level of the venue. In Fig. 4.2**C**, We use the variance in median income across establishments to illustrate: CBDs tend to have lower segregation on aggregate but higher variance between venues, indicating sorting within them. Many of the most segregated CBDs are in cities with even more segregation, and are thus still doing comparably better: as shown in Fig. 4.2**D**, downtowns are typically far to the integrated side of their parent city's distribution.

The consequence of this is that the scaling relationship found in prior work [335], which says that larger cities have greater exposure segregation, does not hold when we isolate central business districts. Fig. 4.3 shows that scaling holds for each type of business we explore here, but not for CBDs. This suggests that much of changes to exposure segregation across cities of different sizes is due to changes at the periphery.

Expected and observed mixing are converging

The pandemic changed decisions around both remote work [53] and social exploration [321]. Given the importance of downtowns to aggregate mixing in cities, the next logical line of inquiry addresses whether cities will become more segregated business districts become less important economically and socially.

To assess changes in segregation patterns, we compared observed segregation values to those predicted by our NULL model. Rather than relying solely on correlation coefficients, which capture only linear relationships, we calculated the mean absolute deviation from the NULL prediction line (where observed equals expected segregation). For

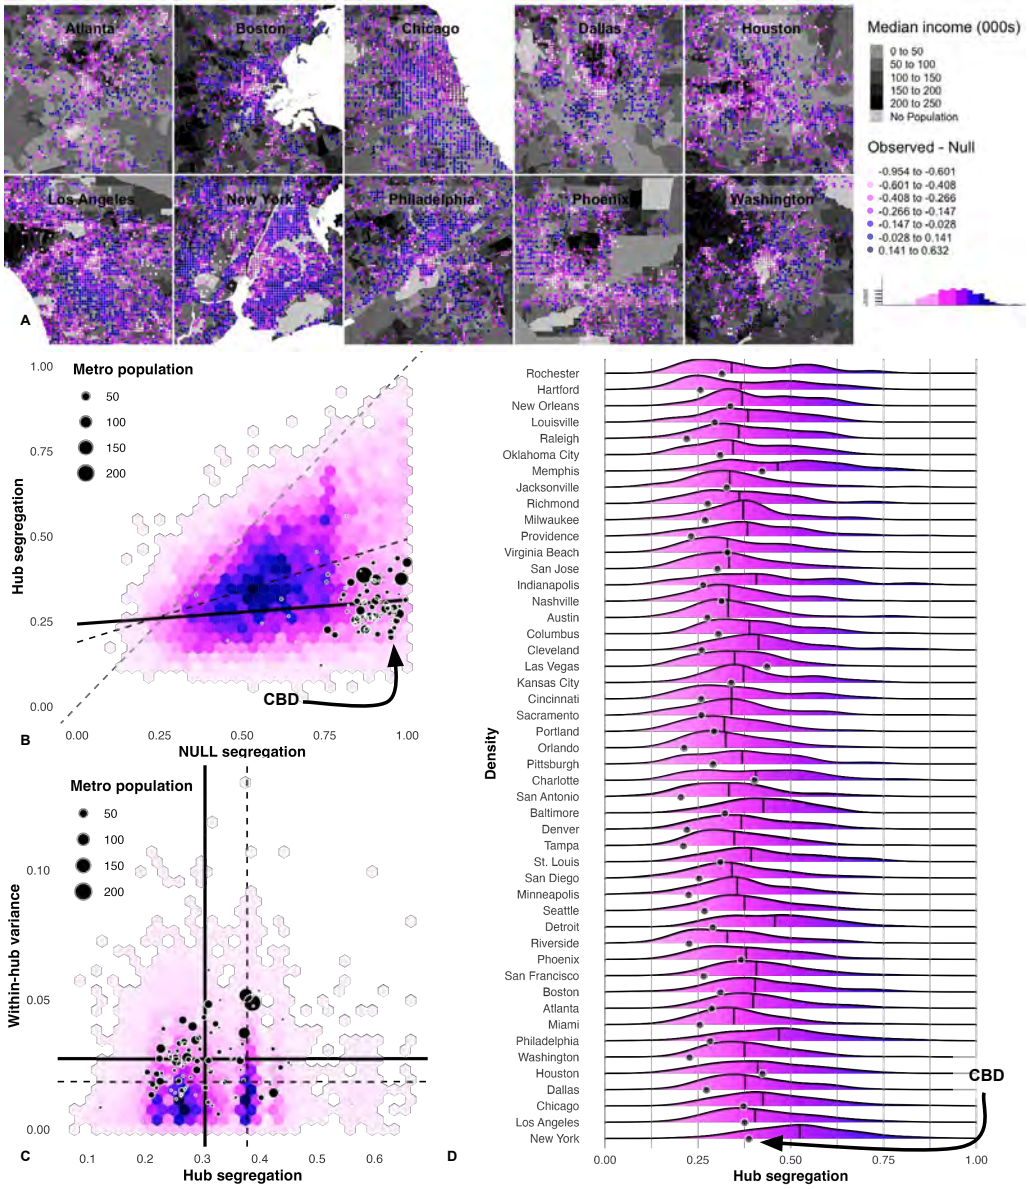

Figure 4.2. **Hubs and mixing.** **A** Maps illustrating the difference between bridging potential and place segregation, showing distinct zones of mixing—usually downtown—and segregation. **B** Plots showing the correlation between bridging and mixing at the level of the hub, with downtown clusters extracted. **C** Heightened mixing in these clusters does not come with greater variance, generally, which would suggest that people are not sorting within clusters. **D** Distributions for all hubs across the top 50 cities by population, with the downtown cluster extracted, shows that these clusters tend to be at the integrated tail of the distribution.

each location-year observation, we computed the absolute difference between observed and NULL-predicted segregation values. We then averaged these deviations within each year and calculated 95% confidence intervals using standard errors. This metric directly quantifies how much actual venue visits deviated from distance-based expectations, with higher values indicating more unexpected mixing patterns. The extremely narrow confidence intervals (± 0.0005) reflect the large sample size and statistical precision of our estimates.

In examining the pandemic's impact on urban segregation patterns, our analysis reveals both the magnitude and pervasiveness of behavioral changes in restaurant patronage. Prior to the pandemic, restaurants exhibited substantial deviation from NULL-

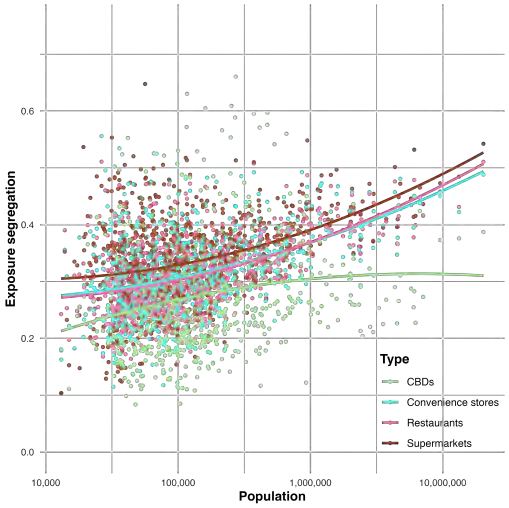

Figure 4.3. **City size and mixing.** On an aggregate level as well as for certain kinds of amenities, exposure segregation scales with city size, but this relationship breaks down when we isolate central business districts.

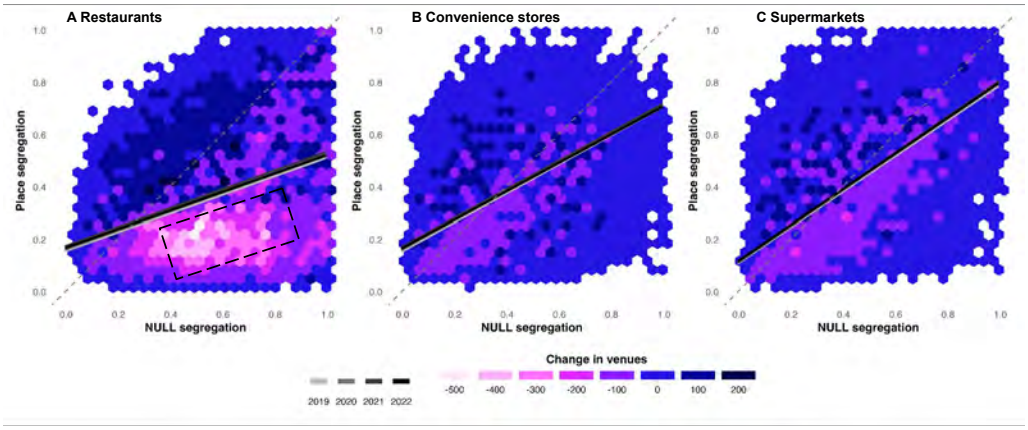

Figure 4.4. **Changes over time.** We plot changes to distribution of mixing and bridging by binning points of interesting according to observed and expected values and examining how the quantities in each bin changes; we also show the linear fit, as above, for mixing and bridging in each year of the data. We can see that for **A** restaurants, the greatest change during the pandemic came from restaurants that were high *expected* segregation but *low* observed segregation, indicating that travel past intervening opportunities, which we see in 2019, declined through the pandemic. We do not see this same dynamic for **B** convenience stores and **C** grocery stores, which did typically had better fit between expected and observed in 2019.

predicted segregation patterns ($\mathbf{0.280}\pm0.0005$), indicating significant cross-neighborhood mixing. The onset of COVID-19 triggered a widespread shift toward more predictable, proximity-based patronage: in 2020, 57% of restaurants (349,480 locations) displayed patterns more closely aligned with NULL model predictions, reducing the average deviation to $\mathbf{0.271} \pm 0.0005$. This trend intensified in 2021, with 58% of establishments (328,102 locations) moving even closer to NULL-predicted values, pushing the average deviation to its lowest point of $\mathbf{0.258} \pm 0.0005$ – representing an 8% shift toward more predictable patronage patterns. The year 2022 marked a partial return to prepandemic behaviors, though notably incomplete: while the average deviation increased to $\mathbf{0.270}\pm0.0005$, only 42% of restaurants (234,097 locations) maintained their pandemic-era patterns, suggesting a gradual restoration of cross-neighborhood mixing. These findings quantify not just the scale of behavioral adaptation during the pandemic, but also the persistent nature of these changes even as formal restrictions eased, highlighting the

potential long-term implications for urban social mixing patterns.

Discussion

In this brief communication we find that the intrinsic mixing potential of an amenity is dependent on its type and its position in the urban hierarchy. Prior work has shown that cities with amenities distributed to “bridge" demographic divides in the population experience more mixing. Yet these results ignore the importance of hierarchy in determining both human mobility and urban structure. Hierarchy is evident in the importance of downtowns, which tend to have low “bridging" but high mixing—for the simple reason that they tend be located in the centre of concentric zones sorted by race, class and age, rather than between them. Downtown offices draw commuters past intervening amenities, generating mixing in a unified central location. This has the strongest effect on restaurants, which are *dispersed* rather than *central* amenities—those for which patrons would prefer not to travel long distances, but end up doing so when that visit is attached to a commute to the office—a *central* place. Supermarkets are not typically part of downtown amenity clusters and for them bridging does indeed predict mixing. Comparing supermarkets stores with convenience stores, we see that the comparably more central supermarkets show a tighter bridging-mixing link than dispersed convenience stores, which suggests that location is more important to the mixing profile of anchor businesses that are attractions in and of themselves than businesses that attach to larger clusters. These results will have particular relevance as many jobs switch to hybrid or remote work, curbing the importance of urban downtowns in the urban hierarchy [380] and placing greater demand on suburban attractions.

Much of this stands in contrast to what we saw earlier: rings of isolation and pockets pockets of segregation limit exposure at the neighbourhood scale, yet central hubs supersede that constraint by attracting diverse visitors past nearer alternatives. These results, focusing on venues instead of broader zones, refine the mesoscopic patterns we saw earlier: geometry predicts aggregate exposure, but hierarchy—who goes downtown, and why—governs mixing within clusters. Many restaurants sit far off the bridging–mixing diagonal because they draft off of agglomerative effects and trip chains; supermarkets hew closer to the diagonal as classic “dispersed goods” [122]: when they are dispersed between communities, they bridge, but when they are dispersed within communities, they isolate–and the fit between expected and observed is tight.

The “central place effect” we identify clarifies an ambiguity in the previous chapter: affluent central neighborhoods may avoid ambient diversity, yet these downtown anchors remain engines of mixing for the city as a whole. In light of rings and pockets that expanded during and after the pandemic, this chapter shows that observed mixing converged toward distance-based expectations during that same extended period, consistent with the retreat from downtowns, a decline central places, and a reduction in trip chains [53].

Thus, meso-scale geometry explains where isolation aggregates; micro-to-macro hierarchy explains when and how mixing persists despite that geometry. Policy levers shift accordingly: pockets need connections; rings need anchors; and central clusters need to remain thick enough to justify skipping intervening opportunities.

The intervening-opportunities NULL model treats all destinations within an industry category as equivalent intervening opportunities—e.g., within “restaurants”, a highway McDonald’s, a mid-range neighborhood bistro, and a Michelin-star venue enter the model with equal weight when counting closer alternatives. This abstraction is useful for isolating the effect of structure, but it can misstate the effective choice set when venues are vertically differentiated (price, quality, cuisine), capacity limited, or serve distinct pop-

ulations. In those settings, the model may misclassify "true substitutes", inflating the intervening-opportunity field and therefore shifting expected segregation. A natural extension would weight opportunities by similarity (price tier, ratings, cuisine embeddings) or estimate category-specific/nested choice sets so that intervening opportunities reflect characteristic differences rather than industry classifications. In the absence of data on this, we note however that this model accounts for a large share of the variation we see in mobility *without* any additional information on quality, suggesting that intervening opportunities are an important determinant of visitation.

Methods

We construct our measure for socio-economic mixing with data from SafeGraph [426], a location services provider, which comes in the form of origin-destination flows from neighborhoods (Census block groups) to points of interest. SafeGraph gathers location data from smart phones by aggregating GPS logs from applications that have obtained user consent to collect such data; the sample constitutes $\sim 10\%$ of the population. They assign visits to points of interest by clustering GPS pings and joining these clusters to adjacent building polygons, using relative distances and time-of-day to manage conflicts [426]. Previous work has shown that SafeGraph data are demographically calibrated [465], and this data underlies the study on mixing and bridging referenced here [335]. We use this rectangular matrix consisting of 220,000 origin home block groups (a Census aggregation with a population ~ 1000) and 7 million destination points of interest like restaurants, grocers, and other businesses. We infer the socio-economic strata of the visitors from each origin with Census estimates from the American Community Survey [498] of median income in that area.

Measuring mixing. We measure mixing by looking at a measure of balance proposed in [321] that asks, in what proportions do visitors from different socio-economic strata visit an amenity? If a restaurant attracts visitors from just one income bracket, we consider that amenity to be segregated; if it attracts visitors from different brackets in equal proportion, then we consider it to be mixed. Following the convention of earlier work, we consider the segregation S of an amenity α to be a distance from an ideal scenario where people from all socio-economic classes visit in equal proportions. This is defined as follows

$$S_\alpha = \frac{5}{8} \sum_q \left| \nu_{q\alpha} - \frac{1}{5} \right|, \tag{4.1}$$

where q represents an income quintile and ν represents the portion of visitors from that quintile. We scale by $\frac{5}{8}$ so that each value spans 0 to 1, with 0 being perfect integration (equal proportions from all classes) and 1 being perfect segregation (visitors from a single class). Each quintile is calibrated to the metropolitan area, rather than the nation as a whole.

Computing bridging. We are interested in generating an expected segregation estimate to compare to the observed. To do this, we model flows to amenities using an intervening opportunities framework, which generates flows from origins to destinations based on both distance and the number of intervening opportunities [469], and then compute the same segregation measure on the simulated data. Classic intervening opportunities models typically use areal units like counties or tracts as origins and destinations; here we

have origin neighborhoods represented by census block groups and destination amenities, stratified by business type (defined by 6-digit NAICS code) so that we are not apportioning visits between, for example, restaurants and grocers. Following the core insight of intervening opportunities theory that the probability of interaction depends on both distance and intervening alternatives, we model how far origin i is from destination j and count the number of comparable (same NAICS) destinations within a circle of radius d_{ij} centered at origin i, rather than attempting to count only destinations that lie precisely between i and j. We define the probability of interaction between origin i and destination j as

$$P_{ij} \;=\; \frac{e^{-\beta d_{ij}}\,(1+r_{ij})^{-1}}{\sum_k e^{-\beta d_{ik}}\,(1+r_{ik})^{-1}}, \tag{4.2}$$

where P_{ij} is the probability of interaction between origin i and destination j, β is the distance-decay parameter, d_{ij} is the great-circle distance between i and j (meters), and r_{ij} is the *intervening-opportunities rank*: the number of same-NAICS destinations strictly closer to i than j (we use $1 + r_{ij}$ to avoid division by zero at the nearest destination). All quantities are computed within NAICS classes: for each origin i and class c, let $J_i^{(c)}$ be the set of destinations with NAICS c; both r_{ij} and the denominator are computed over $k \in J_i^{(c)}$. We obtain $\hat{\beta} = 0.00005$ by maximum likelihood on observed trips (conditioning on each origin–NAICS total). In particular, we estimate this distance–decay parameter β in a destination–choice (conditional logit) formulation: letting n_{ij} denote observed trips, we maximize

$$\mathcal{L}(\beta) = \sum_i \sum_c \sum_{j \in J_i^{(c)}} n_{ij}\, \log P_{ij}(\beta), \tag{4.3}$$

conditioning on each origin–NAICS total $N_i^{(c)} = \sum_{j \in J_i^{(c)}} n_{ij}$. We set $\hat{\beta}$ to the value that maximizes this objective and use it throughout. We then resample the observed trips from each origin–NAICS cell according to $P_{i\cdot}^{(c)}(\hat{\beta})$ and compute the null segregation using Eq. 4.1. Our NULL model ignores differences in quality within broad industry classifications, but it nevertheless serves as informative.

Remote Work and Urban Structure

5 Remote work and urban structure

Abstract In the following chapter, we explore the “doughnut effect"—the idea that COVID-19 and a subsequent rise in telecommuting will create durable gains for suburbs at the expense of downtowns—with a focus on day-to-day activity. Real estate markets are equilibrating in response to changing preferences around living and working, with movers often choosing suburbs and exurbs over urban cores, as recent work documents; less clear is how these changes affect mobility and spending patterns. We address this by documenting 6 years of mobility, from January of 2019 to June of 2024, across 382 metropolitan areas in the United States. We use a set of two-by-two matrices to classify flows and then observe changes to these flows over time. We show a persistent shift in urban activities to the suburban “doughnut": trips from peripheral to central neighborhoods have declined, as have trips from population centers to job centers; instead residents are traveling more within and around the suburbs. This has also curbed flows between predominantly white and predominantly nonwhite communities, potentially reducing intergroup interaction. The strongest drift relative to 2019 has occurred in “superstar” cities like New York and San Francisco. We verify our findings using point-of-sale transaction data to show that many of these superstars are seeing reduced sales downtown—reducing possible sales tax revenue. These findings have important implications for discussions of the urban “doom loop", but they are also relevant to our earlier findings: if activity shifts away from downtown, it will not only reduce integration—as those areas are critical zones of mixing—but it will also redistribute spending from nearby pockets of segregation and into rings of isolation—some of the most affluent parts of American cities.

Introduction

Remote work has altered the demand for different kinds of amenities in cities: larger homes—with room for an office—are now more important; longer, fewer-days-per-week commutes are now more tractable [152]. The consequences of remote or hybrid work arrangements on residential and commercial real estate are becoming clear: a shift from city to suburb in the demand for housing alongside a reduction in demand for downtown offices [500]. Yet the consequences for the economy are deeper than real estate: workers in industries amenable to work-from-home are now present during business hours in neighborhoods that may have emptied out daily before the pandemic. To understand this, we construct a panel of data on spatial interaction spanning 6 years and the contiguous United States, with the goal of describing the changing nature of travel in cities, suburbs and exurbs since 2019—a period which begins with the height of “superstar” cities [218], includes the pandemic and corresponding shocks to mobility, and ends with years of adaptation as preferences around remote work become integrated into the economy. Our focus is on generating a large collection of descriptive statistics to shed light on urban life and its evolution over the past five years. We relate these data to information on remote and hybrid work offerings across the country to show that many of the areas with the strongest adoption of work-from-home have seen the largest changes.

Evidence from repeated, nationally representative surveys shows that work-from-home stabilized at roughly 20–30% of paid days by 2023–2024, indicating a structural

break rather than a transitory shock [52, 53, 162]. A growing body of economic literature addresses the considerable effects of the shift to remote work on real estate [500], including commercial vacancy [217] and residential redistribution [380]. Although the pandemic has abated, both work-from-home [53] and social distancing [54] continue, albeit in evolved forms relative to 2020. Time use data show the emergence of the "introvert economy", with earlier dinner reservations, fewer hours with friends and more on screens [447].

Less clear is how these changes redound to trip patterns. In 2020 and 2021, many studies explored shocks to human mobility and its possible consequences for both viral spread [518] and the economy [82], but these occurred while those shocks were still present. Human mobility has been a focus of urban science for many years, with studies for two decades [98, 205, 101]. Many studies of human mobility to date consider interesting questions about innovation [41], the nature of intergroup interaction [38, 321] and corresponding socio-economic opportunity [299, 132]. Yet most work focuses on narrow time windows. This is because data on human mobility are proprietary, limited, and new: in other words, it is difficult to acquire and collect from private sources, difficult to construct robust samples because of sparsity, and difficult to observe changes over time because the data have only been available for a short period of time. This creates validity concerns, which are compounded by the changing nature of human mobility across the period for which most researchers have access to data; for example, how relevant is 2017 to 2024 given considerable changes to our economic geography over that period? Yet the time window available for researchers to study is growing, and here we address these limitations by building a time series that spans 6 years, from 2019 January to June 2024. We corroborate these shifts with point-of-sale data, which show parallel declines in downtown spend and gains in suburban/exurban areas.

Our focus on trip patterns is motivated by their relevance to infrastructure policy, municipal budgets, and socio-economic mixing. Although few studies explore the persistence of changes induced by COVID-19, the pandemic appears to have had durable effects on consumer behavior in the US; notably, drive-through restaurants had seen activity rebound to prepandemic levels by the end of 2022 while restaurants without drive-through infrastructure had not [315]. Changes like these indicate greater reliance on the automobile, fewer opportunities for spontaneous interactions, and lagging recoveries in business districts. Given the importance of monitoring these aspects of urban life as cities equilibrate to new technologies that enable remote working and reduce commuting pressures, we modify an existing method that collapses complex mobility patterns into a simple matrix to track trends over time.

Our goal is to expand our understanding of these changing preferences by developing a series of mobility signatures, reducing the web of flows between origins and destinations to a series of matrices. Earlier work shows that we can compare cities by reducing the full origin-destination matrix to a 2×2 matrix that measures flows within and between "hotspot" and "non-hotspot" areas [286]. This method classifies trips according where they begin and end: those going from busy or "hot" areas to other hot areas are called integrated I; those starting in quiet or "cold" areas and terminating in other cold areas are called random R; and cold to hot and hot to cold are called convergent C and divergent D, respectively. Our methodological contribution is twofold: we reframe this specific method as a generic operator that can be applied to a richer array of attributes than just activity, like tracking social mixing by partitioning on demography; we then show that its value is greatest as a tool for monitoring long time series. Formally, we treat the classifier as an operator on flow matrices that is parameterised by an ordering variable (e.g., centrality, demography) and fixed 2019 thresholds, so each city–month maps to a

3-simplex (I, C, D, R), or probability vector that sums to 1; we call the resulting trajectories **ICDR-TS**, for time series. Fixing partitions makes the signatures comparable over time and across cities. This generalises the method into a family of signatures with a common geometry and comparable time series, and provides a clear justification for the dimensionality reduction, which necessarily obscures and abstracts away information: it enables us to monitor just four interpretable variables over time.

This yields comparable trajectories that visibly encode regime shifts: 2020 displacement, partial reversion, then a plateau. A growing literature compresses evolving mobility networks into latent factors or clusters—via dynamic stochastic block models [303, 355], matrix/tensor factorisation [258, 169], and manifold/embedding methods [127, 214]—offering strong predictive performance but limited interpretability for policy. Relative to factor or embedding methods, the value in our approach is interpretability and stability: the axes are defined ex ante (activity, centrality) and the shares can be validated against other datasets like point-of-sale transactions.

We contribute to the existing literature on mobility in five ways. (i) Data: we assemble, to our knowledge, the longest multi-city panel of aggregate mobility in the U.S. (2019–2024). (ii) Method: we generalise ICDR into a monitoring operator (**ICDR-TS**) that yields comparable time series across spatial, economic, and demographic partitions [286]. (iii) Measures: we introduce a *doughnut index* for core–periphery reweighting and a *mixing index* for between-group interaction. This mixing index is simpler than what we considered prior—enabling us to monitor long time series—and subtly different because it considers how neighbourhoods interact rather than residents. (iv) Validation: we benchmark changes in activity against point-of-sale spending at 1.1M points of interest to confirm behavioral shifts. (v) Comparative patterns: we document size gradients in reweighting, connecting to established urban scaling regularities [70, 68].

The pandemic catalyzed a durable reallocation of where people live and work, with remote and hybrid arrangements reducing the daily value of proximity to downtown offices [52, 217, 379]. What we still lack is a long-horizon, comparable view of how day-to-day mobility and spending subsequently reorganized within metropolitan areas. In the following, we assemble a panel of mobility activity for 382 cities, which we manage by compressing each month's origin–destination matrix into a set of signatures that are interpretable and comparable across time and space. This lets us quantify the "doughnut effect" in activity, building on existing work in housing and office values. In doing this, we create a lens through which to understand remote work, commuting and supercommuting, the "doughnut effect", downtown vibrancy and the urban "doom loop"—all important issues that will frame policy debates about cities for the foreseeable future.

Results

We gather origin–destination (OD) data for all 382 metropolitan areas in the US from 2019 through 2024. These take the form of spatial interaction networks, which we show in Fig. 5.1**A**. Looking across the US for such a period requires dimensionality reduction. To do this, we borrow and modify an approach that reduces mobility to a two-by-two matrix of high and low values, looking at flows between those quadrants [286]. In its original conception, this involves measuring flows between areal units with different levels of activity: *integrated* flows are those from high-activity origins to high-activity destinations, *convergent* flows are those from low-activity origins to high-activity destinations, *divergent* flows are the opposite, and *random* flows connect low-activity origins to low-activity destinations. While this process can occur with only information contained within the OD matrix, here we assess several other dimensions on which we can divide and classify parts

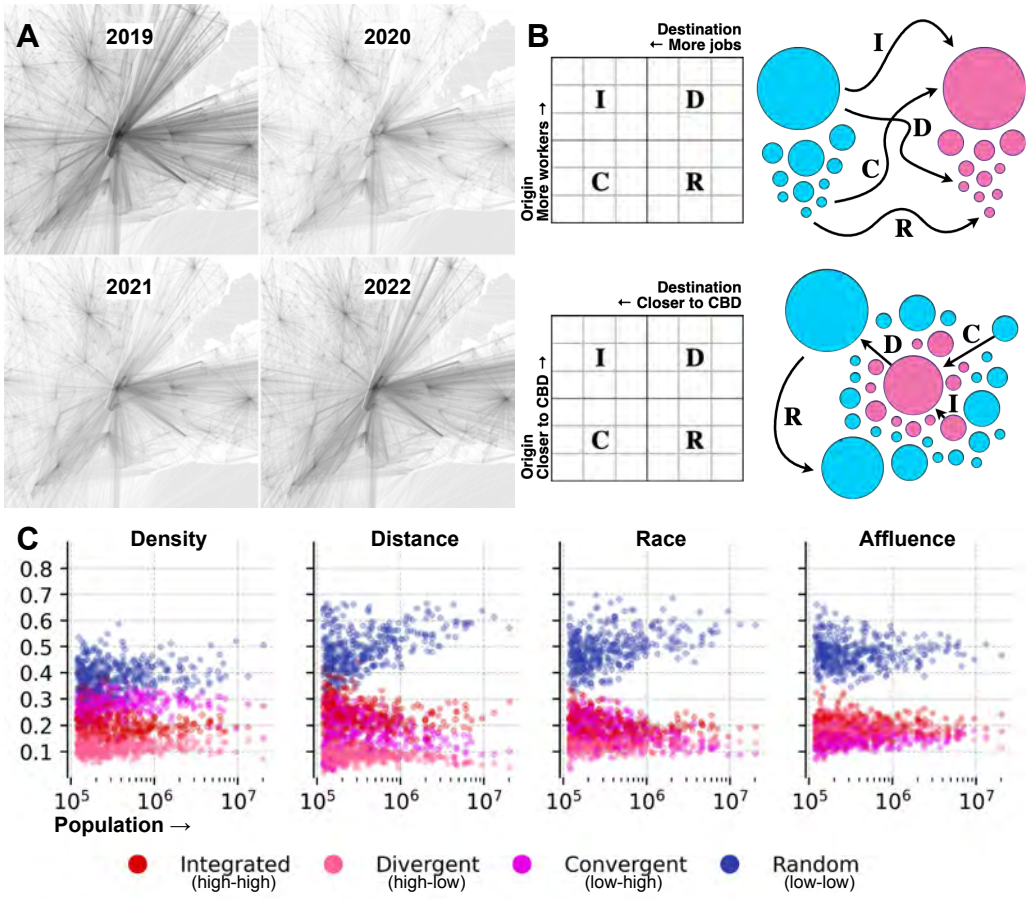

Figure 5.1. **The process.** **A** The origin-destination network for flows in New York City over 4 years of our data, showing a marked drop in mobility as well as an attenuation of activity in the core. **B** A diagrammatic explanation of our method: we sort the flow matrix, then partition it so it becomes 4 aggregates; we can do this for spatial or aspatial attributes, for example sorting by workers at the origin and jobs the destination, or sorting by distance to the central business district for both rows and columns. **C** When we do this for 2019, we can plot all cities according to their ICDR values to compare them; generally, as cities get larger, random flows away from the core become a larger share of all flows—an indication that larger cities are more polycentric.

of a city, including demography and geography. We do this by ranking rows and columns not according to the marginals but according to other characteristics. We call this generalised matrix the **ICDR** matrix, and we demonstrate it in Fig. 5.1**B**. For demography, we use race and education: we divide areas into high-minority and low-minority classes, as well as an alternative specification that uses the fraction of residents with a college degree. For geography, we use distance to the central business district (CBD), which we identify as the largest cluster of businesses obtained via DBSCAN [180] (see Methods for more detail), splitting into central and peripheral. Along with high-population and low-population classes, these give us insight into the changing nature of American economic geography since 2019. We partition origins and destinations into high and low using a fixed 2019 tercile split per city, considering the tracts in the top tercile along the relevant dimension—distance, density, race, and so forth—to be hotspots and those in the bottom two terciles to be non-hotspots. A median split blurs this head–tail structure; terciles isolate important locations while retaining sufficient mass for stable inference. Results are robust, however, to top-half and top-quartile partitions, which we show in Appendix Figs. B.1 and B.2.

This normalization makes magnitudes directly comparable across cities and years.

Note that quantiles keep partitions rank-invariant and cross-city comparable over time. This makes movement in I, C, D, and R represent a real shift in where trips go, rather than an artefact of changing thresholds in a new month. We choose this because cities—especially principal municipalities with taxes and services to consider—care about whether people are reallocating away from places that were central in 2019, not whether an area is "relatively" central in a given month. Also note that ICDR is unitless and normalized to the interval [0,1]; we therefore report changes in 'ICDR units.' A difference of 0.05 corresponds to five percentage points of the full possible range on this scale.

We show the static view for different variables in 2019 in Fig. 5.1**C**, noting that peripheral flows—"random" according to distance—grow with city size, indicating polycentricity and agreeing with previous literature [288]; other kinds of flows show no clear pattern—in contrast to the original work on Spanish metropolitan areas [286].

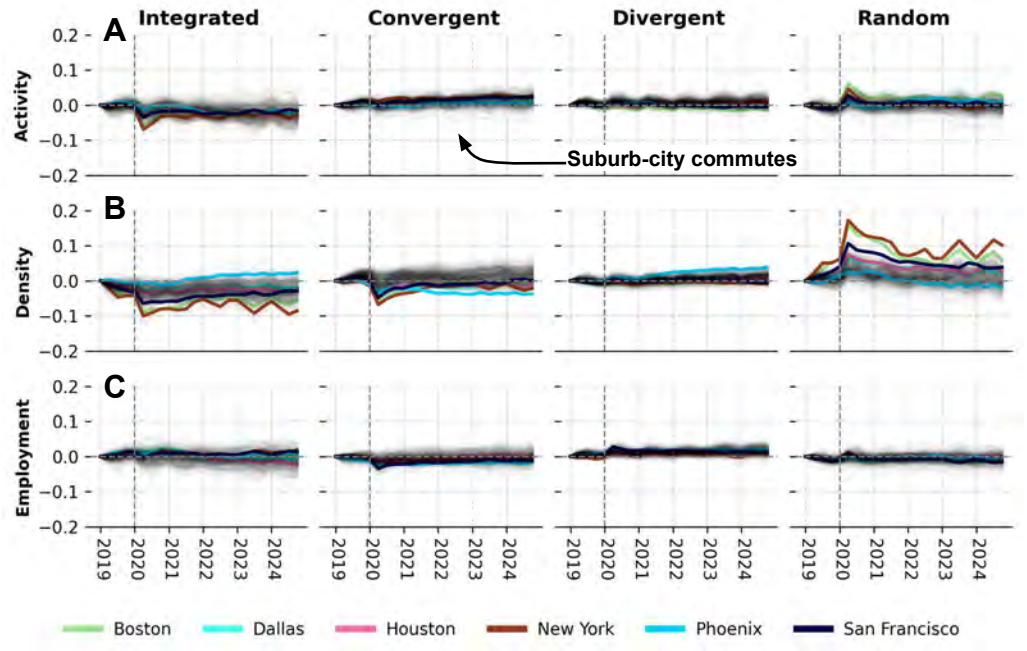

Figure 5.2. **Economic trends.** Quarterly ICDR paths are demeaned to the first quarter of 2019. **A** *Activity* shows small change, although this approach partitioned on marginals from each month, so the hotspots move with the data. **B** *Density* reveals the doughnut: *integrated* flows, dense → dense, fall and remain depressed while *random* flows, sparse → sparse, rise and remain elevated; *convergent* falls and *divergent* stays near flat. **C** *Employment*, workers → jobs, moves less than density, pointing to larger changes in consumption and leisure demand than job relocation.

A doughnut effect in mobility

In all cases, we see a flat trend before April 2020 in Fig. **5.2**, suggesting that mobility patterns stable in the absence of shocks; since 2020, however, the data have been more volatile, indicating that many cities are still adapting to pandemic shocks. In **5.2A**, we use the original approach to partitioning the OD matrix, by computing the marginals—row sums and column sums—on the mobility data, rather than incorporating any administrative data, and sorting on these data to produce hotspots and non-hotspots, indicating limited change. Yet because this approach recomputes hotspots from the OD matrix each month, they can move. To avoid this we consider flows from population-dense to amenity-dense areas in Fig. **5.2B**: integrated flows, those from density to density, fell during the pandemic and have not recovered in most cities; these are flows within cities, and suggest

that COVID restrictions had the most pronounced effects in urban areas—in line with heightened risk from close proximity—but also that these behaviours have not recovered in many cities. Random flows, from sparsity to sparsity, rose and are still above prepandemic levels. Further, trends for convergent and divergent flows indicate that people are traveling less from population-sparse suburban areas to amenity-dense urban areas, while people are traveling more in the other direction—the change is smaller. The increase in random flows and corresponding decrease in convergent flows indicate a clear doughnut effect. Flows from areas with many workers, in Fig. **5.2C**, to areas with many jobs was comparably less affected.

We test if our results are products of our partitioning by using the median and top quartile rather than the top tercile in Appendix **B.1** and **B.2**; the moves are stronger at the more balanced partition, but limited relative to other measures. Generally, although magnitudes vary, after April 2020, integrated flows fell, random flows rose, while divergent and convergent flows were more mixed over time.

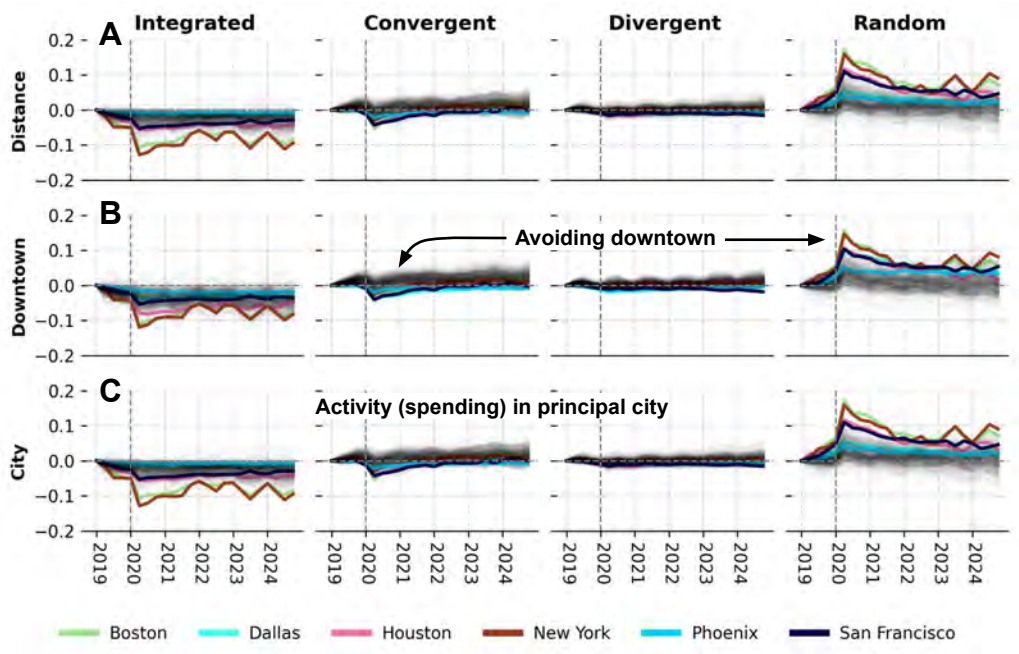

Figure 5.3. **Spatial trends. A** *Distance to CBD* shows the clearest doughnut: *random* periphery → periphery jumps and stays high, *integrated* core → core, declines, and *Convergent* (periphery→core) rises modestly. **B** Considering a downtown obtained by clustering points of interest yields the same signature, reflecting a persistent avoidance of downtown that only partly reverts. **C** Inside the principal city, activity recovers less than in suburbs in many metros, with many municipal governments dependent on sales and business taxes that come from commuters spending downtown.

We explore this further with flows between and within central and peripheral tracts in Fig. **5.3A**. In April 2020, flows between central neighborhoods decreased and flows between peripheral neighborhoods increased. That is, more trips now begin and end in suburban tracts; fewer trips begin and end in urban tracts—evidence for a doughnut effect. Further, after an initial fall in flows from peripheral to central tracts in the first months of the pandemic, these suburban-urban flows rose above prepandemic levels in late 2020 and have exceeded those levels in most cities for the 3 years since. This suggests a change in the relationship between core and periphery, and is robust to alternative specifications for urban/central and suburban/peripheral categories. For example, in **5.3B** we partition according to which tracts are contained in the central business district, defined as the largest cluster of amenities identified with DBSCAN, setting a 500m threshold.

In **5.3C** we identify the principal city in each metro, which has implications for municipal finance since these cities often depend on sales tax revenues to fund services. The trend is strongest in expensive superstar cities [218] like San Francisco and New York. The stability across Dallas and Phoenix is a function of the distributed nature of those two cites: from 2019 to present day, they have the highest rates of random flows and the lowest rates of integrated flows when we consider measures of centrality.

Across these spatial measures, convergent flows drift up and Divergent stay flat—consistent with a "doughnut" pattern rather than a suburbanization of jobs. Using distance ordering, by 2024, nine in ten metros have lower integrated flows than in 2019. Yet at this point in our data, convergent flows are also up on average, indicating that commuting is still important. Because the ICDR matrix sums to 1, this suggests that demand for consumption and leisure activity is shifting to the suburbs.

These are generalities: there are exceptions to each rule and these exceptions tend to occur in smaller cities; larger cities—especially superstar cities that had problems with affordability and housing access prior to 2020—tend to see exaggerated effects. We also see that there is significant divide between "Rustbelt" and "Sunbelt"—colder, older cities in the Northeast and Midwest are experiencing greater change than warmer, newer ones. As we see across Fig. **5.3B**, convergent flows to downtown Phoenix, a sprawling sunbelt city, and random flows within its suburbs changed during the pandemic, it has now stabilized at or near prepandemic levels. In denser rustbelt cities like New York, large shocks have not yet stabilised.

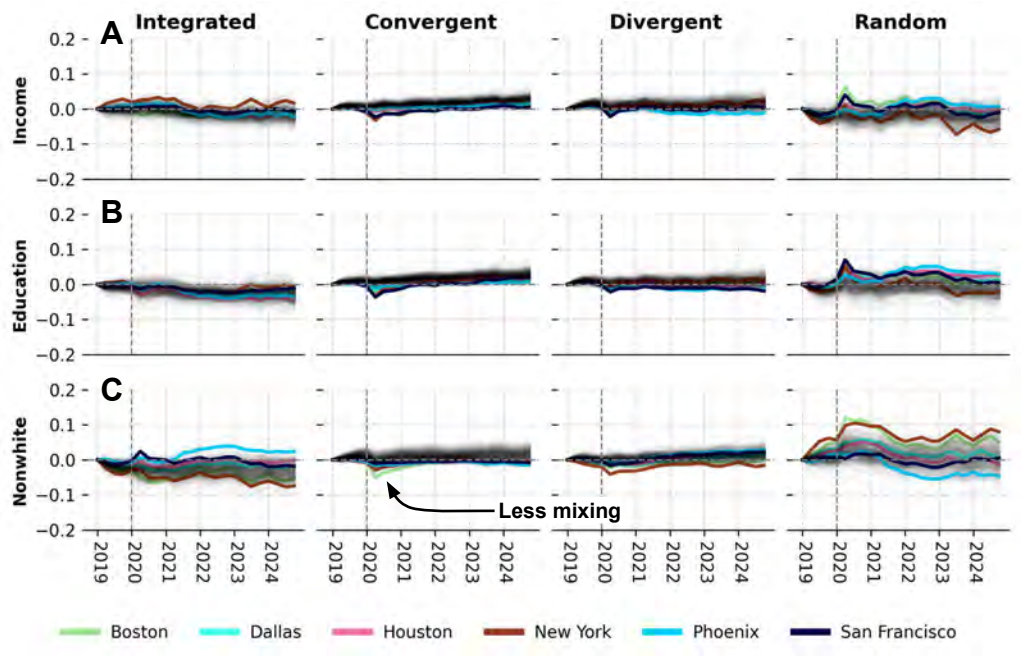

Figure 5.4. **Social trends.** Partitioning by *Income*, *Education*, and *Nonwhite* share shows smaller magnitudes than geography. **A** Income and **B** Education show an initial drop in convergent and divergent flows indicating less socio-economic mixing, but the trend is flat by the end of the study period. **C** Partitioning on nonwhite share, however, *random* low →low strengthened relative to 2019 while cities appear to diverge considerably over time: some cities, like Phoenix, are mixing more, while others, like New York City, are mixing less.

We explore socio-economic mixing in **5.4A** and **B**, with an initial change during the pandemic and many cities recovery thereafter. Cities appear to be dispersing in random—low-education or low-income—flows over time, while convergent and divergent flows are monotonically growing. Yet changes we see in the structure of mobility

do have consequences for racial mixing in cities: relative to 2019, Fig. **5.4C** shows fewer trips crossing from predominantly white to predominantly nonwhite neighborhoods, while the opposite—trips from nonwhite to white neighborhoods—changed little through the pandemic in most cities, possibly because jobs held by members of nonwhite communities were less amenable to remote work [21]. Trips originating and terminating in neighborhoods with similar racial compositions, be they white-white or nonwhite-nonwhite, increased during the pandemic. Although changes in nonwhite ICDR are smaller and noisier than other variables of interest, the increased dispersion in 2024 relative to 2019 indicates that some cities are experiencing less mixing between groups today.

Putting precise numbers to it, on aggregate, mixing is up, however. Sorting by *nonwhite* share, intergroup mixing rises on net, using our mixing index (see methods): we consider the change in mixing to be $\Delta M \equiv (C+D)-(I+R)$, and it shifts by $+\mathbf{0.038}$ in 2024 compared to 2019. Income and education show the same qualitative tilt but smaller magnitudes: 2024 against 2019, $\Delta I = -\mathbf{0.016}$ and $-\mathbf{0.023}$ alongside $\Delta C = +\mathbf{0.017}$ and $+\mathbf{0.019}$, respectively.

Cities are equilibrating at ICDR signatures that represent a hybridisation of prepandemic 2019 and pandemic 2020. Plotting the vector of changes for each city from 2019 to 2020 in **5.5A** and for 2019 to 2024 in **5.5B**, we can see that flows became more random in early on but they have not yet returned to normal. Further, the shock of the pandemic is still reverberating most in the largest cities by population. From 2019 to 2024, the largest 25 metros more than double the shift in $\Delta(R-I)$ relative to the smallest 25, $+0.05$ against $+0.018$.

Because these results are noisy, we construct a doughnut index, representing change in random flows relative to integrated flows, $\Delta(R-I)$. In other words, $\Delta(R-I)$ tells us how much cities have shifted away from core–core travel towards within-periphery travel since 2019. We then regress each city's change in $\Delta(R-I)$ since 2019 on its $\log$ population. Across the *density* ordering, the average change in our doughnut index, is $\overline{\Delta(R-I)}_{2024-2019} = +\mathbf{0.020}$; $\mathbf{72}\%$ of metros raise R and $\mathbf{80}\%$ lower I. The doughnut index in particular has a positive size gradient $\beta_{\log N} = \mathbf{0.007}$ ($p < 0.001$). Using the *distance* ordered, the shift is larger: $\overline{\Delta(R-I)}_{2024-2019} = +\mathbf{0.031}$; $\mathbf{77}\%$ of metros increase R and $\mathbf{90}\%$ decrease I, and the effect steepens with size $\beta_{\log N} = \mathbf{0.014}$ ($p < 0.001$). We report scatter plots of the scaling relationship between change and city size, ordering on distance, in Appendix Fig. **B.3**.

Largest positive $\Delta(R-I)_{2024-2019}$ using distance are rustbelt cities, like Philadelphia $+\mathbf{0.195}$, Washington $+\mathbf{0.156}$, Burlington, $+\mathbf{0.149}$; negative $\Delta(R-I)_{2024-2019}$ values concentrate in smaller sunbelt metros like Safford and $-\mathbf{0.315}$, Flagstaff $-\mathbf{0.122}$), both in Arizona. Generally, larger cities show a stronger shift toward the periphery, a positive size gradient for the doughnut index $\Delta(R-I)$); in particular, ΔI falls more in larger metros. Further, *smaller* cities drive a rise in periphery core flows, a negative size gradient for ΔC).

Two regularities anchor our results: first, distance and density orderings show a persistent reweighting from core $\leftrightarrow$ core and periphery $\leftrightarrow$ core to periphery $\leftrightarrow$ periphery flows; second, the magnitude of that shift scales with city size. We show this in time series in Fig. **5.6A** and **B**. The clearest stratification occurs for convergent flows, which points to fewer commutes in larger cities but more in smaller cities.

On the *distance* ordering, convergent flows rise in most places, with a mean $\Delta C_{2024-2019} = +\mathbf{0.016}$. Although, $\mathbf{88}\%$ of all cities gain, the gains are concentrated amongst smaller cities: the size gradient is negative ($\beta_{\log N} = -\mathbf{0.007}$, $R^2 = 0.284$) and Spearman's rank correlation between population rank and ΔC is $\rho = \mathbf{0.497}$. Moving from the smallest to largest population quartile, the average change steps down from $+\mathbf{0.023}$ to

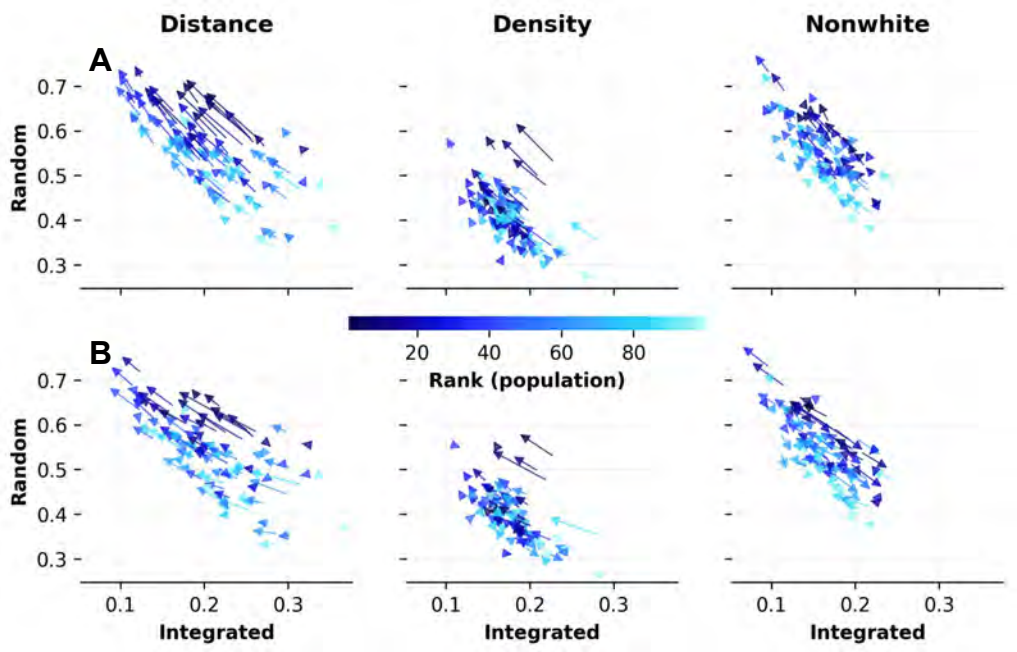

Figure 5.5. **Change vectors.** **A** Change from 2019 to 2020, with a large shift from integrated flows to random flows, especially amongst the largest cities—superstar cities that were becoming expensive prior to the pandemic. **B** From 2020 to 2024, these same large cities are the ones most likely to be far off 2019 values.

$+\mathbf{0.020}$ to $+\mathbf{0.015}$, to $+\mathbf{0.005}$, and the share with $\Delta C > 0$ falls from $\mathbf{98}\%$ to $\mathbf{66}\%$. On the *density* ordering the story repeats, but with a lower magnitude: mean change is $\Delta C_{2024-2019} = +\mathbf{0.005}$ and $\mathbf{61}\%$ of cities gain; the size gradient is again negative ($\beta_{\log N} = -\mathbf{0.008}$, $R^2 = 0.238$), and Spearman's $\rho = \mathbf{0.518}$. By smallest to largest quartiles, the averages run $+\mathbf{0.015}$, $+\mathbf{0.010}$, $+\mathbf{0.002}$, $-\mathbf{0.007}$, and the $\Delta C > 0$ share drops from $\mathbf{83}\%$ to $\mathbf{30}\%$.

Returning to social mixing, we see again that change is contingent on size: since 2019, trips between areas with different demographic compositions—be it race, education or income—rise in most places, and the rise is dominated by smaller cities. The gradient with size is negative ($\beta_{\log N} = -\mathbf{0.009}$, $R^2 = 0.262$); Spearman's rank correlation between population rank and $\Delta(C + D)$ is $\rho = \mathbf{0.502}$, both pointing to stronger re-mixing in smaller metros and a muted response in the largest. That split is stark in levels: across the **25** biggest metros the mean change is essentially flat ($+\mathbf{0.001}$, $\mathbf{44}\%$ positive), while the rest average $+\mathbf{0.022}$ with $\mathbf{91}\%$ positive. Leaders include sunbelt cities like Santa Fe, New Mexico ($+\mathbf{0.089}$), El Centro, California ($+\mathbf{0.056}$), but also smaller cities like Salisbury, Maryland ($+\mathbf{0.064}$), and Parkersburg, West Virginia ($+\mathbf{0.062}$). Declines are concentrated in rustbelt cities like Cleveland, Ohio ($-\mathbf{0.024}$), Detroit, Michigan ($-\mathbf{0.021}$), and Baltimore, Maryland ($-\mathbf{0.020}$).

Computing the change in intergroup mixing with I and R flows as well, such that $\Delta M \equiv \Delta[(C + D) - (I+R)]$ tells the same story in one number: the mean shift is $+\mathbf{0.040}$ with a negative size gradient ($\beta_{\log N} = -\mathbf{0.018}$, $R^2 = 0.249$; Spearman's $\rho = \mathbf{0.506}$), again indicating that smaller metros drove most of the recovery in mixing.

Redistributing spending

The fiscal and economic consequences of the shift are borne out in point-of-sale data, which we use as a validation. We ask whether dollars are spreading across more places

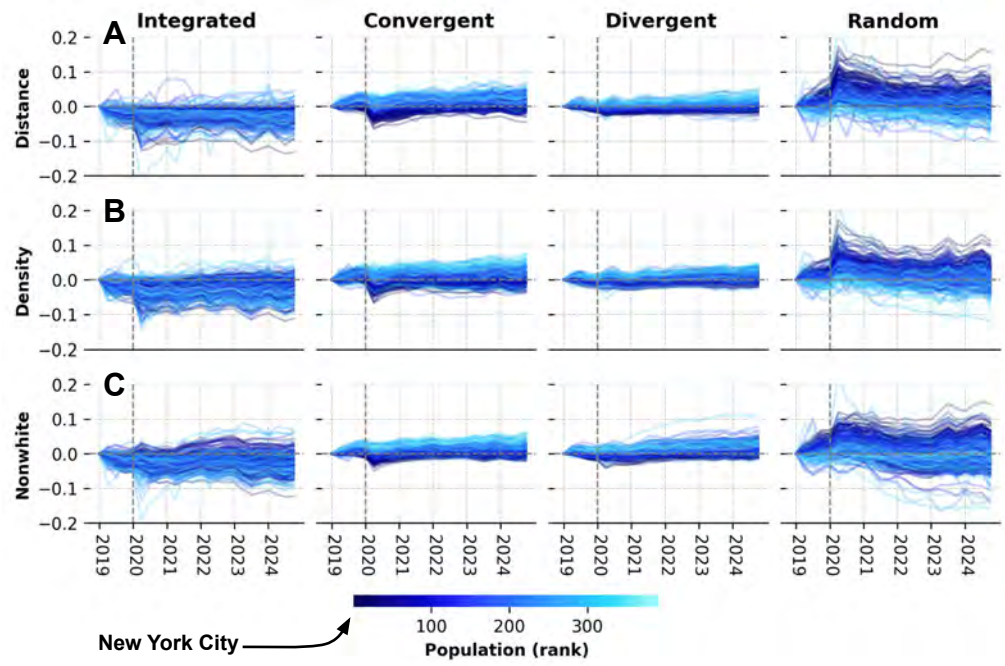

Figure 5.6. **Size and change in mobility composition.** We plot the same ICDR trends as before, colouring metros by population rank to show how size corresponds to change over time. **A** Looking at *distance* to the CBD, we see there is clear stratification of large and small cities for convergent and divergent flows, with steeper declines in the larger metros, indicating a weaker core-periphery dynamic; though the relationship is noisier, we see that random flows gained more in larger metros, indicating a relative shift to towards polycentricity. **B** Looking at density, we see the same pattern holds—persistent elevation of convergent and random flows. **C** Thinking about social mixing by partitioning on nonwhite population, the results are more mixed but many large metros, including New York City, still have greater shares of random, high white population → high white population, flows.

or concentrating into fewer. Rather than diffusing, spending has become *more* unequal since 2019. This indicates that, as we saw in the previous chapter, downtowns actually help redistribute activity—and without them, it concentrates.

At the tract level, concentration rises on both transactions and spend. The pooled Gini for transactions increases from 0.724 in 2019 to 0.773 in 2024; for spend it rises from 0.770 to 0.803. The share for spending in the top decile of tracts moves in parallel: transactions held by the top 10% of tracts climb from 0.569 to 0.641, and spend from 0.638 to 0.693. Dynamics are not monotone: concentration dips early in the pandemic and then reconcentrates, with troughs at 0.708 and 0.764 in April of 2020 for transactions and spending, respectively, and peaks at 0.780 and 0.814 in late 2022. The number of active tracts remains large throughout, at 65,800, so these shifts reflect redistribution rather than a changing coverage.

Taken together, the spend results align with the ICDR–TS patterns: the reweighting away from core ↔ core trips does not imply a broad flattening across many neighborhoods. Instead, activity and dollars reconcentrate into fewer tracts—often peripheral hubs and surviving anchors—consistent with a thicker ring of periphery ↔ periphery travel and thinner core-bound chains. A pooled Lorenz curves and Gini time series appear in Appendix Fig. **B.4**, where we show that this is more a rustbelt than a sunbelt phenomenon, as with the above data: New York City became more unequal while Houston became more equal. We must note, however, that there is not relationship between city size and these changes to the distribution of spending.

This is manifest in a shift towards suburbs and way from downtowns. We aggregate

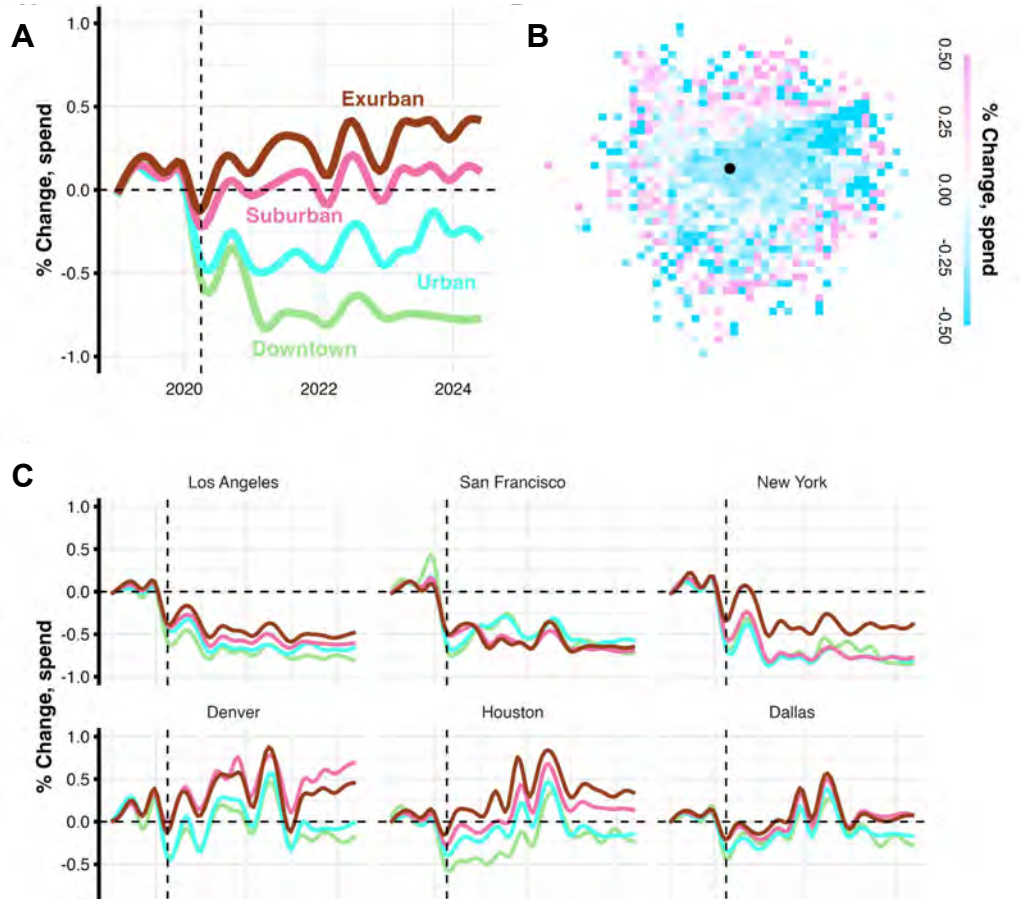

Figure 5.7. **A** Changes in spending according to binned population density, with spending in downtown and urban areas contracting since the pandemic and suburban and exurban areas expanding. **B** A synthetic metropolitan area comprised of the average of all cities in our data shows a clear doughnut, with activity expanding on the periphery and contracting in the center. **C** Disaggregated by city, we see that the results are heterogeneous: in superstar cities that experienced rising costs of living before the pandemic, like Los Angeles, San Francisco and New York, spending has fallen in all categories; less dense, less expensive cities, spending is flat in urban and downtown areas but rising in suburban and exurban areas.

data on spending to the level of the tract and classify those tracts according to population density: following earlier work [380], 99th percentile is downtown, 90th to 99th is urban, 50th to 90th is suburban and the bottom half is exurban. Spending at restaurants downtown has fallen 75% since 2019; at suburban and exurban restaurants, it has risen by 15% and 45%, respectively. Urban restaurants excluding downtowns have seen spending decline by 25%. These data mirror changes to home values [380], suggesting, as we would expect, that people moving to the periphery dine out at the periphery as well, at the expense of the center.

In order to observe an aggregate doughnut effect, we also produce a synthetic metropolitan area by setting downtown of every city in our data to the coordinates 0,0 and, with every city now stacked in the same artificial space, compute the mean change in spending for each 1km^2 grid cell. We see in Fig. 5.7**B** a noisy but prevalent ring over growing expenditure in the periphery while the center experiences a contraction.

When we split the data by city in Fig. 5.7**C**, however, we see that the changes to spending have been heterogeneous. In cities that are defined as superstar [218], those with growing property values and policies limiting new construction, spending has declined across all densities. In San Francisco for instance, as of the end of our study period, spending in cities and suburbs is down by 75%; in New York, spending is down in both

city, suburb and exurb—though the exurbs have only seen a drop of 50%. Spending across Los Angeles and its periphery has also fallen. In cities that were not facing housing crises going into the pandemic, notably Houston, Dallas and Denver, spending is up in suburban and exurban areas and is near prepandemic levels in urban and downtown areas.

Across all cities in 2019, 14% of all tax revenue comes from sales tax and a further 2% comes from business licensing, which could also decline if fewer restaurants and bars locate downtown when fewer patrons commute in each day [279]. Recent data suggest that changes are occurring: in 2021, sales tax revenue fell to 13% of the total. Looking across the past 2 decades in Fig. **5.8A** and **B** shows that these sources of money for cities have been stable for many years and the changes brought about by the pandemic threaten to upend this stability. We saw above that spending is down in Los Angeles, for example, and that city depends on sales tax for 20% of its revenue. San Francisco depends on business licensing, with more than 20% of revenue coming from these fees in recent year; we might see a lag as businesses decided to shut down, but the decline could be similar since both spending and mobility data show that San Francisco is experiencing a strong doughnut effect.

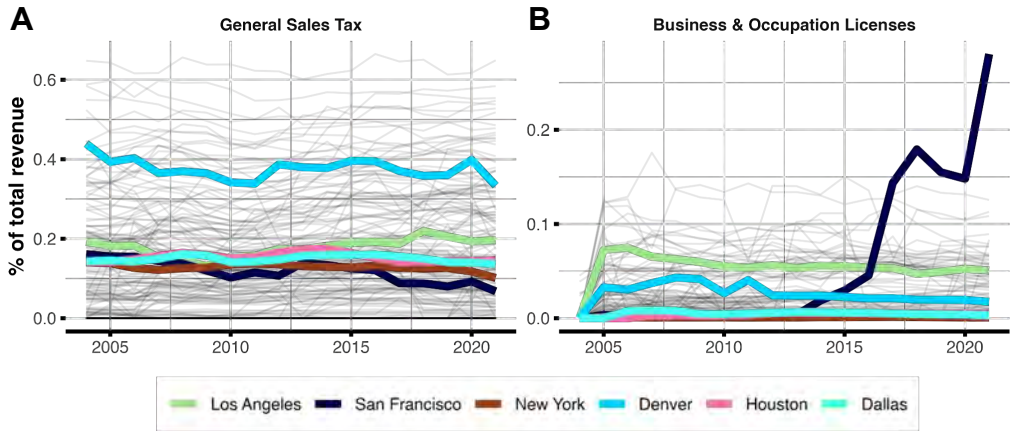

Figure 5.8. **A** Many cities are dependent on sales tax for revenue, especially Denver, which has seen flat spending trends in the city since the pandemic. **B** Although San Francisco is comparably less dependent on sales tax, it does receive much of its revenue from business licensing.

Discussion

In this study we have also introduced a simple yet effective way to monitor changes to mobility across a large sample of cities over a long period of time, extending ideas from previous work [286]. Our method prioritises description over prediction to enable policy evaluation. We see that the adaptations set in motion by the pandemic—most notably remote work—are still present in cities, which, while intuitive, have consequences for both policy and how we interpret early work on urban mobility. Cities have settled into a hybrid of 2019 and 2020: integrated, central movement remains depressed, random, peripheral movement remains elevated, and the cities that most resemble themselves prepandemic are those that are smaller, sparser, and cheaper.

On the distance ordering, the average doughnut index $\Delta(R-I)$ increased by $+0.031$ and intensifies with city size $\beta_{\log N} \approx 0.014$, consistent with a stronger shift toward polycentric travel in bigger metros. Results are similar for density ordering, but more mixed for measure that sort on education, income and race—although even here we see a clear distinction between large and small cities.

Spending follows foot traffic. Point-of-sale data show persistent contraction downtown with compensating gains in suburbs and exurbs—especially in superstar metros—implying pressure on tax bases, municipal services, downtown businesses. A particular risk is feedback: fewer commuters, fewer amenities, and a less vibrant downtown thus drawing even fewer commuters. Although real estate constitutes the lion's share of municipal budgets, sales taxes are also important—and foot traffic has subtler effects on urban living, from "eyes on the street" deterring crime to spontaneous interactions generating innovation. Because production often drives consumption in cities, with commuters eating breakfast, lunch or dinner within the confines of business district, a large redistribution of journeys from core to periphery will also redistribute expenditures, possibly more equitably [285], but with adverse effects on municipal budgets. With productive workers isolated from each other, research suggests that rates of innovation may fall [517, 320].

Yet the same feedback can, in principle, run in reverse: a periphery that becomes thick with third places, street life, and reliable connections can become a mixing surface rather than a set of parallel lives. These locations will not minimise the mean free path length as a focal point would, so we might still expect mixing to decline; however, there is an opportunity to create more sustainable urbanism that still generates planned and unplanned encounters through places that exist between the lowest and highest layers of central-place hierarchies: the proximal neighbourhood and the central business district.

How the new direction of travel moves with or against the existing urban infrastructure may constrain changes. Infrastructure will encourage or discourage certain kinds mobility: dominant radial highways facilitate travel in and out of the urban core; abundant orbital highways allow for point-to-point journeys around the suburban periphery. We can think of the change we see in the data as a transition from the radial to the orbital city, and its implications are manifold.

Our findings agree with models predicting residential decentralization and altered commuting when work-from-home rises, while employment need not decentralize at the same rate [159]. Using the ICDR-TS operator we propose here, future work can embed in difference-in-differences or event-study designs that leverage work-from-home exposure by industry mix or firm policy to move from description to inference. It also offers a tractable target for structural calibration (matching $\Delta(R-I)$ and $\Delta(C+D)$ by city size) to discipline counterfactuals around transit investments, office conversions, and zoning changes.

A strength of our study is the fixed quantiles, which allow for comparison over time and space, but this is also a limitation. An important trend, documented in the literature, is migration away from urban cores during the pandemic [380]. Our partitions do not relabel areas that truly do shift rank—emerging suburban and exurban satellites. Capturing such structural reranking would require periodic rebasing to population, which falls outside the scope of this paper given the cadence of official releases.[1]

Our work also has implications for research on human mobility, especially regularities and patterns: at least in the US, travel preferences are equilibrating: data from the last year show that ICDR signatures are stabilising, even in cities that have not returned to prepandemic norms. Prior to 2020, mobility research was beginning to reveal important insights about urban life [321, 40] and human behavior [18, 444]; for example, research using data from 2017 suggests that amenities situated between rich and poor areas foster socio-economic mixing [335], and in the previous chapter we showed that downtown is a critical area for the aggregate amount of mixing because it defies this "bridging" model—attracting diverse groups at higher rates than the model predicts. Whether or not

[1] Although, we partially mitigate this by reporting robustness with alternative centre definitions (CBD clusters and distance) and with median/quartile splits.

similar work using data from 2022, 2023 or 2024 would reliably convey the same ideas about cities is an open question. Our work suggests that we are entering a period of stability and thus that new research, leverage recent or current data, would provide information about the new regime that we are in postpandemic. We provide an initial compendium of results and ideas here: commuting patterns have changed, integration—according to both race and class—have changed, and spending choices have changed. Superstar cities in particular have entered a new regime.

Data and methods

We use mobility data from Advan, a location intelligence company that derives foot traffic by aggregating GPS pings that occur within building polygons. Our data take the form of an origin-destination matrix where the origin is always home, rather than an intermediate stop, and the destination is a point of interest. Advan sources GPS data from ~ 35 million devices, isolating clusters of multiple GPS traces over multiple minutes that fall within a given building [6]. Advan data capture a representative sample of the US population with small biases (± 0.05) by race, age and gender [276]. To reduce bias, we aggregate these data from the Census block group to the tract, and use months rather than weeks as our periods. We consider only leisure trips—which we class as those to grocery and retail stores, restaurants and bars, cafes and bakeries, along with museums, arenas and stadiums. This reduces—though by no means eliminates—double counting, if a person, for example, goes to the office and then to lunch. With these data, we construct matrices for each city for each month from January 2019 to June 2024, which to our knowledge is the longest panel of mobility data assembled for academic research to date.

This comes with limitations: although we are able to see the evolution in aggregate mobility over the long run, the devices that constitute the panel will change as users opt in and out of location tracking on certain mobile applications. This leaves open the possibility that biases appear in the data at different times, causing fluctuations that are not the product of changing mobility but of a changing sample. To ameliorate this risk, in another specification we compute our measures with normalised flows, such that the trips emanating from each origin sum to 1—so that we consider the shares rather than the absolute values.

Flow classification. To monitor and describe many cities over a long period, we extend a method for reducing a flow matrix to 4 dimensions [286]. This method considers a matrix F with rows corresponding to origins i and columns to destinations j. It then marginalizes the total out-flows and in-flows for origins and destinations, respectively, and then reorders those rows and columns according to these totals so that in one corner we have the cell with both the most trips emanating from it and the most trips terminating in it, and in the opposite corner we have the cell with the fewest of both totals. Dividing the ordered rows and columns in 2 segments each with a quantile break [414] creates 4 quadrants that represent *integrated* (high-high), *convergent* (low-high), *divergent* (high-low), and *random* (low-low) flows. This matrix is normalized so that $I+C+D+R=1$.

Methodologically, we extend this **ICDR** matrix (integrated, convergent, divergent, random flows) from a static classifier into a *time series monitoring operator*. For each city-month pair c, t, we map the OD matrix $F_{c,t}$ to $(I_{c,t}, C_{c,t}, D_{c,t}, R_{c,t})$, yielding **ICDR trajectories** in a fixed 3-simplex ($I+C+D+R=1$). We generalise the partition beyond activity hotspots to geographic, demographic, economic attributes, preserving the same four quadrants in each case. This design keeps the geometry constant, making signatures comparable over time and across cities, and it foregrounds interpretability for policy by

allowing us to monitor core–periphery dynamics and intergroup connections. We illustrate this in Fig. 5.1**B** and explain in detail below.

ICDR

Let $F_{c,t} \in \mathbb{R}_{+}^{n_c \times m_c}$ be an origin-destination (OD) matrix for city c in month t, with rows i as origin tracts and columns j as destination tracts.

1. **Choose an ordering variable X for rows and Y for columns:**

   - Examples: at origins, choose $X \in \{\textit{density}, \textit{employees}\}$ so that X_i is population density or employed workers in tract i; at destinations, choose $Y \in \{\textit{amenities}, \textit{jobs}\}$ so that Y_j is amenity density or jobs in tract j; distance to CBD $X = Y = \textit{distance}$; demography $X = Y \in \{\textit{nonwhite}, \textit{income}, \textit{graduates}\}$.

2. **Fix partitions to a 2019 baseline.** For each city c and variable X, compute the top-tercile threshold $\tau^X_{c,67}$ on 2019 levels (similarly $\tau^Y_{c,67}$ for columns). Define the "hot" set and its complement

   $$H^X_c = \{i : X_i \geq \tau^X_{c,67}\}, \qquad L^X_c = \{i : X_i < \tau^X_{c,67}\},$$

   and analogously for columns (H^Y_c, L^Y_c), where $F_{c,t}(i,j)$ is the count of trips from origin tract $i \in \{1, \ldots, n_c\}$ to destination tract $j \in \{1, \ldots, m_c\}$ in city c during month t; X_i and Y_j are row/column attributes (e.g., population density at i, or amenity density at j) evaluated on a fixed 2019 baseline;[2] $\tau^X_{c,67}$ and $\tau^Y_{c,67}$ denote the city-specific 67th-percentile (top tercile) cutoffs in 2019.

3. **Compute shares.** Let $T_{c,t} = \sum_{i,j} F_{c,t}(i,j)$.

   $$I_{c,t} = \frac{\sum_{i \in H^X_c} \sum_{j \in H^Y_c} F_{c,t}(i,j)}{T_{c,t}}, \qquad C_{c,t} = \frac{\sum_{i \in L^X_c} \sum_{j \in H^Y_c} F_{c,t}(i,j)}{T_{c,t}},$$

   $$D_{c,t} = \frac{\sum_{i \in H^X_c} \sum_{j \in L^Y_c} F_{c,t}(i,j)}{T_{c,t}}, \qquad R_{c,t} = 1 - I_{c,t} - C_{c,t} - D_{c,t}.$$

   By construction $I_{c,t}, C_{c,t}, D_{c,t}, R_{c,t} \in [0,1]$ and $I_{c,t} + C_{c,t} + D_{c,t} + R_{c,t} = 1$. Further, $\mathbf{s}_{c,t} \equiv (I_{c,t}, C_{c,t}, D_{c,t}, R_{c,t}) \in \Delta^3$ (a 3–simplex). The mapping is scale–invariant in $F_{c,t}$ (multiplying all entries by a constant leaves shares unchanged) and rank–invariant over time because thresholds are fixed to 2019. For the distance specification, I corresponds to core $\leftrightarrow$ core, C to periphery $\rightarrow$ core, D to core $\rightarrow$ periphery, and R to periphery $\leftrightarrow$ periphery. For demographic specifications ($X{=}Y{=}Z$), $C{+}D$ summarizes between-group travel while $I{+}R$ summarizes within–group travel.

We report results for density and distance orderings as standard "doughnut" measures, and for nonwhite, income, education for social mixing. By way of example, in the distance ordering, integrated flows are central $\leftrightarrow$ central and random flows are peripheral $\leftrightarrow$ peripheral, convergent flows are peripheral $\rightarrow$ central (commutes), and divergent flows are central $\rightarrow$ peripheral. We include alternative specifications with downtown set as the largest amenity cluster (DBSCAN [180] with $\varepsilon{=}500$ m) and the principal municipality, defined as the largest city in the metro area. Finally, we partition on the marginals of the OD matrix itself (original ICDR approach) as well as another with *workers*$_i$ and

[2] Fixing thresholds to 2019 keeps partitions invariant through time.

jobs$_j$. Robustness checks use median and quartile splits. Further, to reduce sensitivity to device sampling, we compute

$$\tilde{F}_{c,t}(i,j) = \frac{F_{c,t}(i,j)}{\sum_j F_{c,t}(i,j)},$$

shares per origin. We then reweight these shares by origin population,

$$\hat{F}_{c,t}(i,j) \equiv \tilde{F}_{c,t}(i,j) \cdot \text{pop}_i,$$

so that each tract's contribution is proportional to residents rather than total visitation; ICDR is recomputed on $\hat{F}_{c,t}$.

ICDR-TS

1. **Temporal aggregation.** Convert monthly shares to quarterly with simple averaging to smooth high–frequency noise.

2. **Baseline demeaning.** Define $\Delta I_{c,t} = I_{c,t} - I_{c,t_0}$ etc., with $t_0 = 2019\text{Q1}$. We plot $(\Delta I, \Delta C, \Delta D, \Delta R)$ per city and summarize by population–rank bins (Fig. 22–26).

In particular, we partition using row and column terciles, such that the "hotspot" is the top tercile along each dimension. We choose a tercile rather than a median because many aspects of city life—including the activities we measure—produce skewed distributions: a small share of tracts concentrates population, jobs, amenities, and trips. A median split flattens this. Terciles keep enough mass in each cell of the 2×2 to give stable estimates across hundreds of cities and months; quantiles keep partitions rank-invariant and cross-city comparable.

Modelling

We define the **doughnut index** for city c and time t as

$$\Delta(R-I)_{c,t} \equiv (R_{c,t} - I_{c,t}) - (R_{c,t_0} - I_{c,t_0}),$$

where positive values indicate reweighting from core $\leftrightarrow$ core to periphery $\leftrightarrow$ periphery relative to 2019. We report distributions and size gradients.

When ordering by a demographic attribute Z (e.g., nonwhite share), we define a **mixing index** as

$$M_{c,t} \equiv (C_{c,t} + D_{c,t}) - (I_{c,t} + R_{c,t}).$$

and calculating $\Delta M_{c,t} \equiv M_{c,t} - M_{c,t_0}$ so that positive values indicate more intergroup interaction (relative to 2019) and negative values indicate less, in a single number.

Size gradients

We read the "size gradient" by regressing changes within each city on $\log N_c$: for each metro c,

$$P_c = \alpha + \beta_{\log N} \log N_c + \varepsilon_c,$$

where P_c is the change in the index (e.g., $\Delta(R-I)$ or ΔC) and N_c is population. The coefficient $\beta_{\log N}$ is the size gradient: a one–log-point increase in population (an e-fold) is associated with $\beta_{\log N}$ more reweighting; per doubling, the implied effect is $\beta_{\log N} \cdot \ln 2$. We also report Spearman rank correlations to capture monotone size patterns.

Spend data. We confirm our results and quantify their consequences using point-of-sale transactions from SafeGraph, which monitor $\sim$1.1 million POIs and $\sim$200 million transactions per month [427]. The panel spans 2019-2024. We aggregate transactions to tract, classify tracts into density bins (downtown, urban, suburban, exurban) using the same 2019 thresholds as in ICDR, and compute spend indices normalised to 2019 averages. Specifically, for bin b and city c,

$$S_{c,b,t} \equiv \frac{\sum_{i \in b} \mathrm{Spend}_{i,t}}{\frac{1}{|T_0|} \sum_{t \in T_0} \sum_{i \in b} \mathrm{Spend}_{i,t}}, \quad T_0 = \{\text{2019 months}\}.$$

We focus on consumer-facing categories (e.g., eating and drinking, retail) and smooth monthly volatility with 4-month rolling average, then compare $\Delta S_{c,b,t}$ to ICDR-TS distance/density shifts at the same aggregation. Prior validations show that aggregate spend tracks public company earnings and that unique customers scale with state populations [163, 232].

6 Polycentricity and structure

Abstract Our last chapter showed evidence that cities—and big cities in particular—are entering a new regime, with activity shifting from the core to the periphery. A central debate in urban theory compares the relative advantages of monocentric agglomerations and polycentric conurbations. In monocentric models, business districts concentrate jobs and the matching of employees to employers, but these hubs bring congestion and other costs, forcing jobs and amenities to distribute throughout the city. Modern cities are characterised by tension between centralising and decentralising forces, yet our understanding of how residents actually navigate behavioural demands to produce hybrid structures remains limited. Human mobility patterns offer a powerful lens to understand cities' dynamic structures, revealing complexities hidden in static population maps. Here we analyse the mobility of 15 million residents across 226 cities in France, Germany, and the United Kingdom. We develop a method using mixture models to decompose trips into distinctive scales–proximal, medial, and distal. We find that the lengths of proximal trips—the shortest ones—are stable across the city and between cities while medial and distal distances vary systematically by location and city. Larger cities fragment into ancillary hubs while smaller cities are more dependent on a single core. The primary contribution of this chapter is to introduce a method for modelling and monitoring polycentricity in light of evolving mobility; the results presented are preliminary and illustrative, but the method itself can be applied to work on urban structure going forward.

Introduction

With most of the world's economic production and innovation concentrated in its urban areas [202], understanding and addressing the factors that contribute to the relative function or dysfunction of cities is of vital importance to continued progress and prosperity. Although the structure of cities emerges from spatial and temporal constraints [122, 295], evidence suggests that these emergent forms and functions have economic consequences for life in cities: dense and compact cities are more productive than sprawling and misshapen ones [222], and the residents of connected cities are better off economically [25]. Because the putative mechanisms behind these differences in performance involve reductions to the cost of both communication and transportation [171], it follows that the spatial signatures that define productive and innovative cities manifest behaviourally and socially. The following paper develops a method for capturing the structure of cities at scale with GPS mobility data and shows that variations in structure correspond with important differences in lived experiences. Recent advances in the collection of passive GPS data now allow us to observe signatures for entire systems of cities, at spatial and temporal resolutions that were not possible a decade ago [49]. Our contribution is to bridge theoretical models of urban structure with empirical data on human mobility, using a sample of cities an order of magnitude larger than any previous studies, while providing evidence for the drivers and consequences of this structure.

Cities are not monolithic entities—they are multifarious systems where different activities may thrive under different spatial structures. Central place theory [122] posits a

nested hierarchy of centres—clusters or agglomerations—with different industry mixes; some goods and services must be accessible at short distance, accessed frequently, while other markets can be farther from people because they are required only infrequently. Spatial [528, 295, 435] and social [18] constraints create fundamental limits on how many places we can visit in a given period, which creates an emergent spatial arrangement as businesses, populations, and possibly to a lesser extent infrastructure, adjust to meet these needs. Recent work suggests that these rule hold not just between cities in larger systems of cities, but within cities, as the human need to operate across scales in characteristic ways—close and frequent, far and infrequent—extends across orders of magnitude [17, 444]. For example, restaurants naturally cluster together to produce dining hubs throughout a city [270]. This, in turn, influences the population distribution as people will sort *between* cities according to economic opportunities [138] but *within* cities according to amenities [195].

Urban theory debates the relative advantages of concentrated, monocentric agglomeration and diffuse, polycentric conurbation [373, 67]. In the monocentric model, a central business district concentrates jobs, a thick and liquid labour market where workers and firms and match and rematch without burdensome changes to residential or commercial location [27, 129]. This unified market also enables returns to scale by reducing the need for infrastructure while improving the ease of communication [194]. However, a centralised city also incurs congestion costs [8]: as cities grow, transport networks slow [137] and land rents grow [130] near the core; decentralisation becomes a necessity beyond a certain metropolitan size, leading to the emergence of specialised hubs and clusters throughout the metropolitan area [26]. This transition from monocentric to polycentric is enabled by new forms of transportation [59], and constrained by geographic and demographic factors [429]. The modern city is thus characterised by a tension between centralising forces and decentralising forces [287]. That tension has sharpened evolving as hybrid and fully remote work norms decouple employment from the obligation to be physically close to a single core [162]; understanding its repercussions is therefore urgent.

In contrast to many office jobs, which benefit from close proximity [32], consumption and leisure activities tend to distribute across multiple hubs rather than one singular core [442]. Recent empirical studies of human mobility reinforce this idea: individuals structure their movement within a city at multiple scales, maintaining distinct "activity spaces" for home, work, and other purposes [205, 17, 443]. For example, high-resolution mobility data reveal natural divisions between neighbourhood-centric routines and city-wide trips, implying that the optimal urban form might depend on function. Work trips might still gravitate toward a primary centre to maximise job matching, whereas shopping, socialising, and recreation often manifest in a polycentric pattern of secondary centres in residential areas. Because a nucleus solves the matching problem for jobs, a city may be industrially monocentric yet polycentric for errands and social life, reflecting the varied demands of its distributed functions [17, 444].

The structure of cities has typically been learned from static data—either from remote sensing [112] or from administrative data [306]. Detailed GPS data are transforming our ability to analyse these patterns. Early work on polycentricity has been limited to data on transit trips in a narrow set of cities [423, 95], and studies exploring a broader range of mobility are often contextually limited—one country or a few cities [286, 283, 520]. Bridging stylised urban models with empirical evidence across different contexts and scales requires methods that learn spatial structure without imposing it *a priori* [57]. Here we propose such a method.

Human mobility data offer a powerful lens to understand urban structure, revealing complexities beyond what we can see in population or land use. Tracking where and

when people travel—their daily activities across the city—yields a map of functional urban connectivity, and the "revealed preference" of residents using and experiencing the city. A growing literature uses human mobility data to model and understand behaviour [350], yet these same data can shed light on the external boundaries [165] and internal structures of cities [283, 286]. Mobility patterns allow us to measure the degree to which a city operates as a whole or as a collection of parts; because social connections are important determinants of economic opportunity [211] and these relationships are in turn facilitated by spatial accessibility [46], integration functions as a valuable proxy for economic opportunity. By and large, spatial proximity begets social connectivity [45].

Here we exploit GPS mobility data for urban residents in cities across Germany, France and the United Kingdom to develop a method for understanding how cities balance the forces of centralisation and decentralisation. Our primary goal is to model how residents in different parts of a city allocate time and travel budgets to meet demands for work and leisure. A secondary goal is to then use our model to understand the possible social and economic consequences of different urban structures. We extend existing work inferring polycentric structure from transit trips [95] by applying it to mobility data covering more than 200 European cities. These data capture a rich set of behaviours, allowing us to decompose urban mobility across the spectrum of activities in order to measure how spatial structure varies across functions. Are cities monocentric for the core industries that benefit most from agglomeration, yet polycentric for shopping, leisure, and other idiosyncratic needs? The following analysis tests this proposition with mixture models that decompose the distribution of trips from any given part of a city, which represent a variety of distinct purposes, into a set of component distributions. This mixture-of-scales view makes polycentricity legible as a budget allocation—proximal routines, medial errands, distal commutes—rather than a binary label that a city either earns or fails. In doing so, we find regularities in how people distribute activities in space with implications for social mixing, productivity, and sustainability.

Data and methods

Data

In this chapter, we will begin with data and methods we are demonstrating a method more so than results. We use anonymised GPS mobility traces from $\sim$ 20 million mobile devices across all 226 functional urban areas in France, Germany, and the United Kingdom. Unlike transit smart card data that captures only public transport usage, mobile device data provide insights into a rich variety of trips, including those closer to home that do not require transit and those farther from home that require a car, revealing the complete spectrum of urban mobility. For each device, we identified the home location as the modal `H3` level 10 cell occupied between midnight and 6:00 AM. We applied the infostop algorithm [35] to detect stationary periods and extract discrete visits from continuous GPS trajectories, where a visit constitutes a spatially clustered set of GPS points persisting for at least 5 minutes within a `H3` level 10 cell. Each detected visit generates a trip of length d_{ij}, defined as the haversine distance between the home origin i and the visited destination j. Note that this means that each trip is anchored to home, and that trip chains are not considered as such. This gives us a set of $\sim$ 200 million trips across all cities; London has the most trips in our data at $\sim$ 50 million, and it is the largest city, with a population of $\sim$ 12 million in the entire functional urban area. These trips constitute a spatial interaction network, which we illustrate for London, Paris and Berlin in Fig. **6.1a-c**.

We then use these trips to measure polycentricity with the following approach. We group each device according to its home location, using H3 level 7 cells to aggregate devices so that each model has sufficient data, which we will call a neighbourhood. For each neighbourhood, we construct a vector distances, or lengths L, representing the length of the trip, discarding all other information. The empirical distribution $\mathbf{L}_i = \{l_{i1}, l_{i2}, \ldots, l_{iN_i}\}$ is multifaceted: it aggregates errands near home together with longer commutes and shorter outings. Following earlier work [95], we capture this heterogeneity with finite mixture models of interpretable *scales*. If a city were truly monocentric, meaning that all goods and services agglomerate in a single hub, every trip would approximate the length of the distance between the origin neighbourhood and the nucleus of the city—approximate because even in a monocentric city, amenities will still distribute throughout the area of the nucleus. Under this null, the expected trip length would equal the radial separation between the origin neighbourhood and the nucleus, so that $E[L_i] = r_i$ for all i.

Models

The empirical lengths exhibit pronounced multimodality, suggesting the superposition of distinct mobility processes operating at different spatial scales. To decompose these patterns, we model the distribution of trip lengths from each home hexagon using a finite mixture of Gamma distributions

$$f_i(l) = \sum_{c \in \{p,m,d\}} \pi_{ic}\, \Gamma\big(l \,\big|\, \alpha_{ic}, \beta_{ic}\big), \tag{1}$$

where the mixing weights satisfy $\sum_c \pi_{ic} = 1$ and

$$\Gamma(l \,|\, \alpha, \beta) = \frac{\beta^{\alpha}}{\Gamma(\alpha)}\, l^{\alpha-1}\, e^{-\beta l}.$$

is the Gamma probability-density function with shape α and rate β. The three components $c \in \{p, m, d\}$ correspond to proximal, medial and distal mobility scales. Because trip lengths are positive and usually skewed, we need a kernel whose support is $[0, \infty)$ and whose shape can flex from bell to exponential. The Gamma kernel nests the exponential ($\alpha = 1$) and approximates log-normal for larger α, making it flexible yet tractable.

Our mixture has three components ($c \in \{p, m, d\}$) representing proximal, medial, and distal mobility scales. The result is a reduced description of human mobility, turning spatial interaction networks into six numbers per neighbourhood (π_{ic}, μ_{ic}). The choice of three components balances model complexity with interpretability: three components minimise Bayesian information criterion [329] in 90% of origins, striking a balance between fit and parsimony.

To decompose our data we use the Expectation-Maximisation (EM) algorithm, that begins by examining each neighbourhood's trip length distribution, searching for evidence of multiple processes operating at different scales. Shown in Appendix Fig. **B.5**, a neighbourhood's trip distribution often appears lumpy or multimodal—perhaps showing a concentration of short trips (to local parks and restaurants), another cluster at medium distances (to nearby shopping districts), and a concentration of repeated commutes to the urban nucleus. Rather than forcing this complex pattern into a single distribution, the mixture model asks: might this be multiple distinct distributions superimposed? The EM algorithm discovers these latent components by gradually refining its guess about which trips belong to which mobility process. Starting from an initial partition, the algorithm recognises that some trips are ambiguous but attempts to cluster. Through iteration, it

learns the characteristic distance and spread of each latent mobility type for that specific neighbourhood, decomposing the lumpy distribution into a set of constituent distributions that better explain the observed pattern.

The expectation is initialised with k-means on $\log \mathbf{L}_i$. Logging stabilises the wide span of distances, while k-means gives plausible starting groups for "near", "mid", and "far" trips. The resulting partitions yield the initial parameter set $\boldsymbol{\theta}_i^{(0)}$. From here, the approach is iterative, alternating between E and M steps.

The **E-step**: Given current parameters we compute the *responsibility*—the probability that trip j from cell i with length l_{ij} is from component c, specified as

$$\gamma_{ijc} = \frac{\pi_{ic}\,\Gamma\big(l_{ij} \mid \alpha_{ic}, \beta_{ic}\big)}{\sum_{c'} \pi_{ic'}\,\Gamma\big(l_{ij} \mid \alpha_{ic'}, \beta_{ic'}\big)}. \tag{2}$$

where the numerator measures "how well component c explains this trip"; the denominator normalises across the three components so the probabilities sum to one.

The **M-step**: Treating γ_{ijc} as fractional counts, we maximise the expected complete-data log-likelihood

$$\begin{aligned}
\pi_{ic}^{\text{new}} &= \frac{1}{N_i}\sum_j \gamma_{ijc}, \\
\bar{l}_{ic} &= \frac{\sum_j \gamma_{ijc}\, l_{ij}}{\sum_j \gamma_{ijc}}, \qquad \overline{\ln l}_{ic} = \frac{\sum_j \gamma_{ijc}\, \ln l_{ij}}{\sum_j \gamma_{ijc}}, \\
s_{ic} &= \ln \bar{l}_{ic} \;-\; \overline{\ln l}_{ic}, \\
\alpha_{ic}^{\text{new}} &\text{ solves } h(\alpha_{ic}^{\text{new}}) = -\,s_{ic}, \quad \text{with } h(\alpha) = \psi(\alpha) - \ln \alpha, \\
\beta_{ic}^{\text{new}} &= \frac{\alpha_{ic}^{\text{new}}}{\bar{l}_{ic}}.
\end{aligned} \tag{3}$$

where N_i is the number of trips from neighbourhood i and ψ^{-1} is the inverse digamma function. Intuitively, the first line re-weights each component by its fractional membership; the second and third lines update the Gamma shape (α) and rate (β) so that the new mean and log-moment match the responsibility-weighted sample. The E–M loop repeats until the log-likelihood gain falls below 10^{-6}.

After convergence, we order components by increasing mean $\hat{\mu}_{ic} = \alpha_{ic}/\beta_{ic}$, assigning labels proximal (p), medial (m), and distal (d) accordingly. This ordering ensures consistent interpretation across neighbourhoods: proximal trips represent local movement within immediate surroundings, medial trips capture intermediate movement to nearby activity centres, and distal trips encompass global movement across the urban system.

Note that this process requires three components, and orders those components without any other information. This means that central locations in particular will have a distal component, by necessity, and that the flows comprising this component will represent trips *away* from the centre—because trips within the centre will necessarily be shorter than trips away from it for these locations. This is a limitation of the approach, but we address it in our analysis by examining the weights (travel away from the centre will likely have lower weights if it is occasional leisure) and other characteristics of the component, like the points of interest represented within its flows.

We then link our $\hat{\mu}$ values to a structural variable, r_i, so that the entire spatial organisation of a metropolis can be visualised and analysed more easily. Mixture means allow us to test the assumption that neighbourhood are dependent on the nucleus component-by-component. For each city we identify an *effective centre* as the centroid of the hexagon with the most visits in the city; the radial distance from this centre to hexagon i is r_i. We then estimate

$$\hat{\mu}_{ic} = a_c + b_c\, r_i + \varepsilon_{ic}, \tag{4}$$

with weighted least squares (N_i as weights). Slopes b_c measure how sharply each component or scale responds to urban density gradients. In a polycentric configuration we expect $b_p \approx 0$, as neighbourhood activities remain as such, at all scales, but $b_m > 0$ as residents seek out central places to meet different needs for goods and services, and $b_d > b_m$, as residents commute into the nucleus for work and related activities, like meals near the office.

Extensions

A valuable feature of this approach is that we can classify trips according to these posterior probabilities, and from there we can infer the characteristics of those classes. We leave this for future work, but demonstrate the method here. In order to understand both the drivers and consequences of polycentricity, we are able to append characteristics of trips from each component in each city. For each neighbourhood, we collect the posterior probability distributions for each component ($c \in \{p, m, d\}$), which allow us to assign each trip to its most likely component according to its length. Posterior responsibilities

$$\gamma_{ijc} = \frac{\pi_{ic}\, \Gamma\big(l_{ij} \mid \alpha_{ic}, \beta_{ic}\big)}{\sum_{c'} \pi_{ic'}\, \Gamma\big(l_{ij} \mid \alpha_{ic'}, \beta_{ic'}\big)}. \tag{5}$$

classify every observed trip into the scale most compatible with its length. These assignments enable a suite of downstream analyses—from origin and destination classification to analysis of experienced segregation, as we can now label the origins and destinations according to neighbourhood attributes. A possible focus is social mixing. We are able to extract a rich set of attributes for all cities in our data using remote sensing, available from open sources. For each journey, these include the population of the originated cell, along with building heights and volumes, an index of vegetation, and the venues or amenities available in the terminated cell.

Results

We show the relationship between location and travel in London, Paris and Berlin in Fig. **6.2a**, showing the distribution of journey lengths as the origin neighbourhood moves away from the nucleus, H_1. We can see that most trips are short, and in Paris and Berlin there is also a concentration of trips that follow the line of identity, indicating that they are to the nucleus; this is not present in London, likely the result of London's various nuclei. We set the nucleus in London to Piccadilly Circus; however the city also has Knightsbridge, Westminster, Canary Wharf, and the Square Mile. It is still interesting that Paris does not have the same dispersion as London, given that it also has a new financial centre in La Défense. We can model these distributions with a single component $k = 1$ before moving into our preferred specification in Fig. **6.2b**, which shows the best fit line over component means $\hat{\mu}_i$ for each city. This represents an expected length $E[L_i]$ given a location relative to the centre; these means do not follow the identity line because these

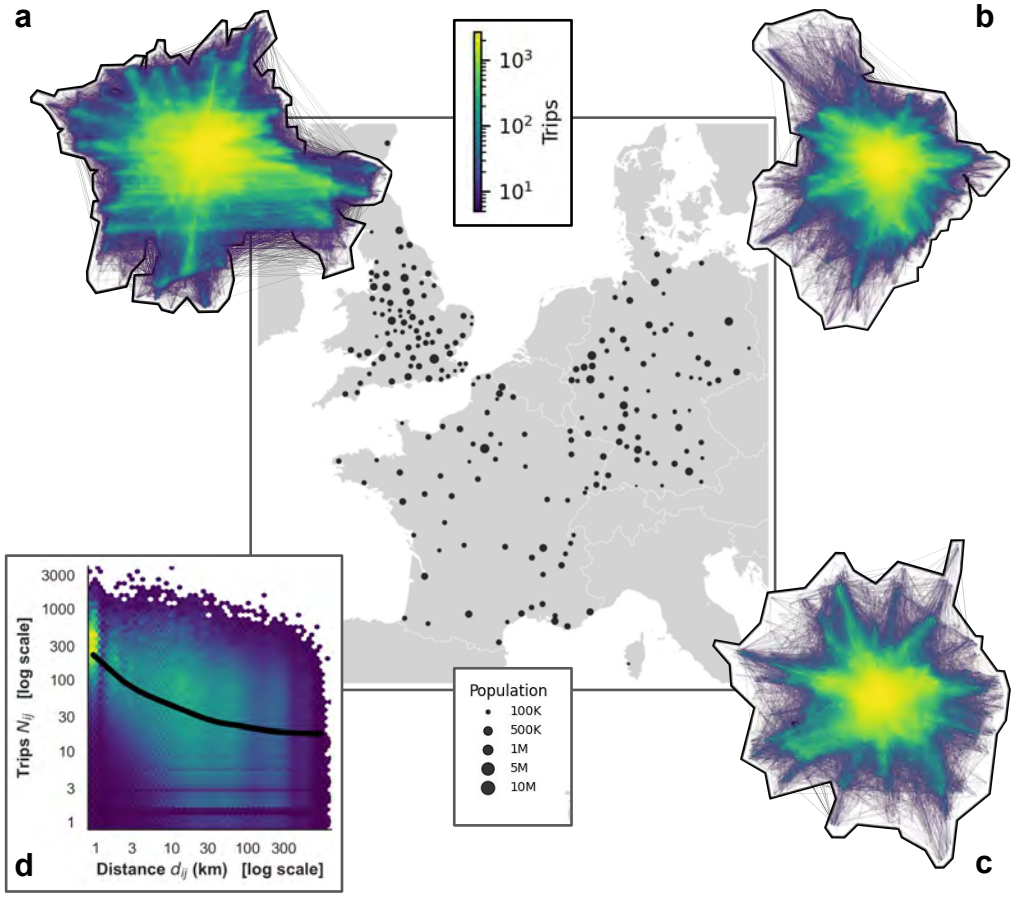

Figure 6.1. **Describing the data.** In the central map, we mark all 226 urban areas in our study, spanning towns of $\sim 10^5$ inhabitants to cities exceeding 10^7. Example spatial interaction networks for **(a)** London, **(b)** Paris and **(c)** Berlin; there is strong core-periphery structure in all networks, but in London in particular smaller hubs are also visible. **d** Aggregate distance–decay in the data for London, showing distance d_{ij} and the corresponding journey count N_{ij}; the curve is a smoothed average of the binned data, highlighting the gradual decrease in trip intensity with increasing distance.

means group short, intermediate and long trips in a single distribution. The inadequacy of these simple models is why we attempt to decompose each neighbourhood's distribution using a mixture of distributions.

We demonstrate our mixture approach with $k = 3$ components on London, Paris and Berlin in Fig. **6.3a**, where we can see that as a neighbourhood gets farther away from the centre—which in London is near Piccadilly Circus—the distal mean $\hat{\mu}^d$ grows, suggesting that these communities need to travel farther to access important services or jobs. The proximal $\hat{\mu}^p$ distances do not systematically vary by distance to the centre, suggesting that local activity operates at a similar scale no matter the neighbourhood, and medial $\hat{\mu}^m$ distances scale but less so than $\hat{\mu}^d$. Paris and Berlin have similar medial weights, as indicated by mix of points from both cities in Fig. **6.3b**. London is often called a collection of villages and we see here that its neighbourhoods typically have higher medial weights, standing out from Paris and Berlin, which is indicative of a city where more needs are satisfied without the nucleus.

We document an interesting fact in **6.3c**: because our modeling approach is agnostic to what kinds of trips it decomposes into constituent distributions, it detects a distal component with high $\hat{\mu}^d$ in central locations. This occurs because our process finds three distributions of trip lengths and orders them to characterise them as proximal, medial and distal; it is agnostic to where those trips terminate. In London, residents inside the

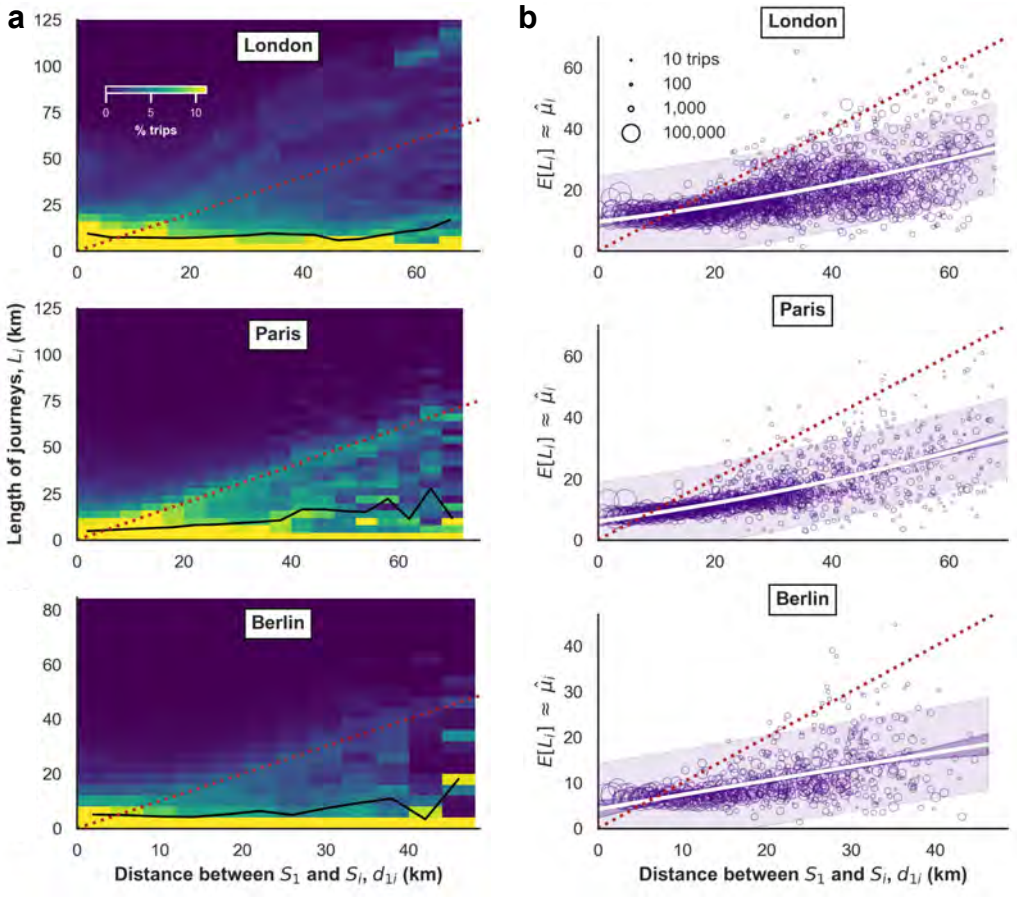

Figure 6.2. **Mobility as a function of relative centrality. a** The relationship between where a trip originates, relative to busiest point in the city, and its length, for London, Paris and Berlin: most trips are short, but we also see longer trips and in Paris and Berlin these long trips concentrate along the identity line, suggesting that they terminate to the nucleus. London is more diffuse, which suggests that it is comparatively less monocentric. **b** For a given home, we compute the expected length of a trip: because most trips are short, this value is less than the distance between the nucleus. Our strategy is to decompose the distribution that gives rise to this mean value for each home.

core travel outside of it and our approach classifies this as distal. These trips constitute smaller proportions for their neighbourhoods relative to suburban ones, but it suggests urban residents might gain flexibility through centrality.

London is the only city of the three where distal trips are longer for residents of the core. In Appendix Fig. **B.6** we show that in these central neighbourhoods, distal trips also have lower weights, despite higher distances—indicating that trips venturing away from the nucleus are still rare. We also see in Appendix Fig. **B.6** that both inner London and inner Paris stand out as places with high medial weights $\hat{\pi}_i^m$ and low distal weights $\hat{\pi}_i^d$—with clear boundaries along the Inner Ring Road in London, which demarcates the original congestion charge zone (now expanded) and the périphérique in Paris. This indicates that neighbourhoods within the confines of the city have better access to amenities and work and have efficient travel distances as a consequence of this accessibility. The same pattern exists in Berlin but it is less pronounced. We also note that proximal components show no relationship between distance d_i and weight $\hat{\pi}_i^p$.

We plot the residuals for the models in Appendix Fig. **B.7**, which captures model fit. While the area in inner London has poor distal fit to go along with low weight, there is also poor fit in an area just beyond the M25 in Slough, which suggests that areas beyond the city limits behave in ways that are not well described by these mixture models. We

do not see comparable areas with poor model fit in Paris and Berlin, however. Variation in model fit comports with London's structure as a collection of villages, wherein these villages have both local services and offices but some dependence on the nucleus: in distant neighbourhoods, many residents will have long trips in the nucleus, which will raise $\hat{\mu}_i^d$, but others will stay local—and this dispersion will degrade model fit. In other words, poor model fit can also signal polycentricity.

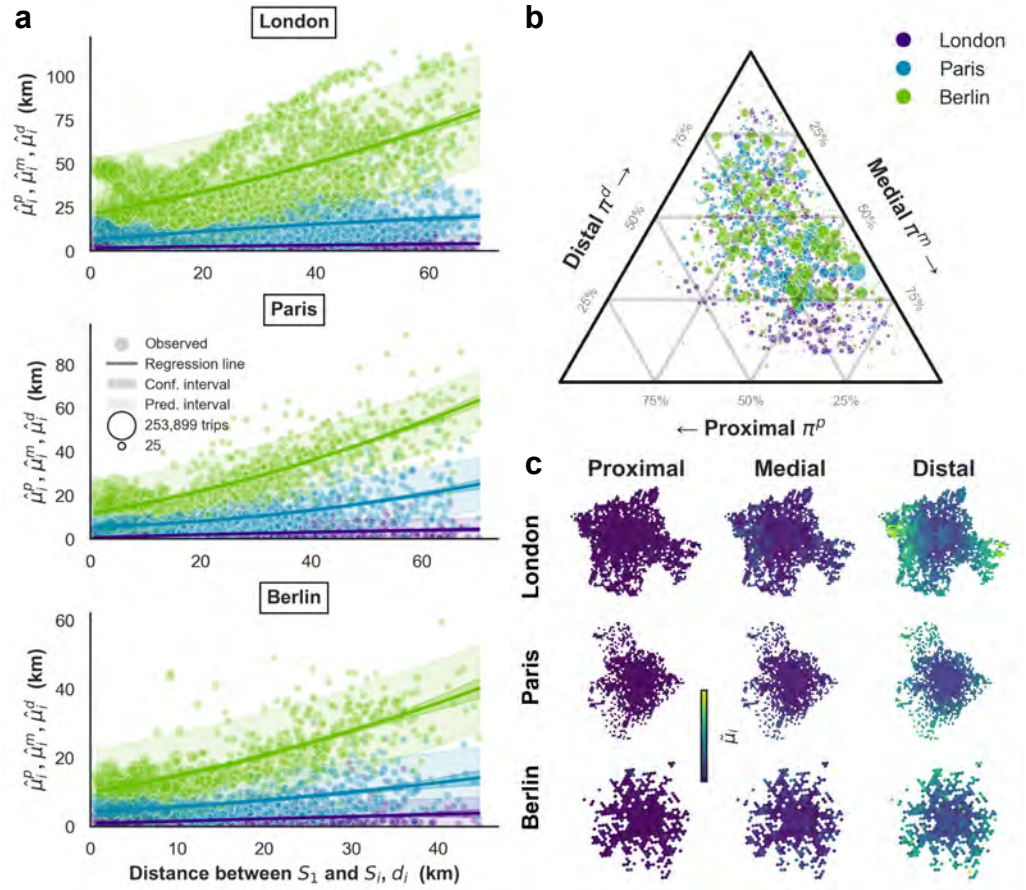

Figure 6.3. **Components of mobility.** **a** Scatter for London, Paris and Berlin, showing proximal, medial and distal component means against distance to the city centre; the proximal component is flat, indicating that people typically stay close to home for many trips, and the medial component grows, indicating that in sprawling suburbs, even local trips become extended. Indicating the importance of an urban core, we see that the distal component grows with distance to the nucleus. **b** Relationship between proximal, medial, and distal weights, or shares of all trips from a given neighbourhood; the proximal component is typically smallest, suggesting that even in urban areas, many still travel to amenities beyond the immediate neighbourhood. **c** Maps of the same three component means across our three sample cities, showing that the most central locations have low means across components while peripheral locations have high distal means. We also see that in London, the distal component represents trips *away from* the core. In these areas, the distal weights are low, even if the distances are high, suggesting that most trips are still local and these trips are different in nature than distal trips in other parts of the city.

Next, we fit mixture models to $\sim$ 20,000 neighbourhoods across $\sim$ 200 cities in Europe. When we normalise each city so that the maximum journey length is 1, and the maximum distance from the nucleus is 1, we can compare all cities to each other. In Fig. **6.4a** we show the same scatters of $\hat{\mu}_i^c$ per component as above, with all cities in our data normalised to produce $\tilde{\mu}_i^c$, and then aggregated. We see a number of important regularities, shown in Fig. **6.4a**: across all cities, as in London, Paris and Berlin, $\tilde{\mu}^p$ is typically location invariant ($\sigma_{\mu^p} \approx 0$) while $\tilde{\mu}^m$ and $\tilde{\mu}^d$ vary more ($\sigma_{\mu^d} > \sigma_{\mu^m} > 0$). The insets in Fig. **6.4a** show the various best fit lines for each city, with brighter colors

corresponding to larger cities; these become steeper from proximal to medial to distal. A steeper line indicates that a city is dependent on the nucleus for those trips. We also see a clear stratification where larger cities have mixed higher and lower slopes but lower intercepts. We see that $\tilde{\mu}^d$ shows a strong positive association with location, but smaller cities, with darker lines, also show a positive relationship for $\tilde{\mu}^m$. This indicates that smaller cities are more dependent on their urban centres while larger cities fragment into subcentres, but the relationship is noisy so we revisit below.

We compute total distance budgets D_i in Fig. **6.4b** by using the weighted sum of components such that $D_i = \hat{\pi}_i^p \tilde{\mu}_i^p + \hat{\pi}_i^m \tilde{\mu}_i^m + \hat{\pi}_i^d \tilde{\mu}_i^d$. We show the simple fact that the cumulative sum of $\hat{\mu}^c$ grows with distance: the peripheries are less efficient than the centres. In **c** we show that the distal component $\tilde{\mu}^d$ scales with distance at a faster rate than the medial component $\tilde{\mu}^m$, which indicates that many forms of travel can remain local even as commuting distance grows; because $\tilde{\mu}^m$ still grows with distance, however, all but the most local trips still demand longer rides farther from the nucleus.

In accordance with research on fixed travel budgets, this implies that speed changes in the periphery as people use automobiles. Yet it also implies a trade-off: if cities are to leverage the nucleus as a place of mixing, as we discuss in the previous chapter, journey lengths will necessarily rise. There will be some ways to economise, since medial trips do not grow in length as fast as distal ones, but distances will rise on aggregate.

The weights of each component are stable across city size and country, which we show in Fig. **6.5a** and **b**. The proximal component is typically the smallest share of trips, while medial and distal components are larger: most needs are met away from home. We also note that most variation occurs along the medial axis of the ternion plot. We can also look systemically at the relationship between city size and the model parameters. In accordance with the monocentric hypothesis, if trips from the distal component focus on the nucleus, b_1 of $\tilde{\mu}_i^d$ approaches 1.

Estimating models as $\tilde{\mu}_{ic} = a_c + b_c \tilde{r}_i + \varepsilon_{ic}$, where $\tilde{r}$ is the normalised radius, or distance from the centre, we see large variations in slope of the regression line with city size in Fig. **6.5c**, indicating a mix of mono- and polycentricity. Within cities, slopes order as $b_p \approx 0 < b_m < b_d$. Across cities, distal slopes b_d are heterogeneous but show no robust monotonic relationship with population, whereas intercepts a_c decline with population across components.

While for slopes, we observe wide dispersion across cities, the dominant size signal sits in the intercepts. In Fig. **6.5d**, the intercept declines with population across components. A higher intercept, a_c, means that even central residents in smaller cities traverse a larger fraction of the urban extent with medial and distal trips than central residents in larger cities. When a high distal intercept coincides with a flatter distal slope, b_d, distal trips are more diffuse and less targeted toward the nucleus: there is a heightened level of distal movement to dispersed destinations within the city even at $\tilde{r} = 0$. This implies that in smaller cities, central residents cover more of the urban extent than counterparts in larger cities. In other words, smaller cities exhibit *generalisation* while larger cities exhibit *specialisation*: residents in big cities concentrate activity on a smaller set of destinations, whereas residents in small cities visit a larger share of the city.

These values a_d and b_d give us two ways to think about a city: nucleus pull (slope), and nucleus diffusivity (intercept) that represent how specialised and targeted mobility, relative to the extent of the city. It appears that nucleus diffusivity scales with the inverse of city size, but nucleus pull does not. Given research that indicates fragmentation city size—cities becoming more polycentric as they grow [26]—these lack of clear relationship here is a puzzle worthy of future work.

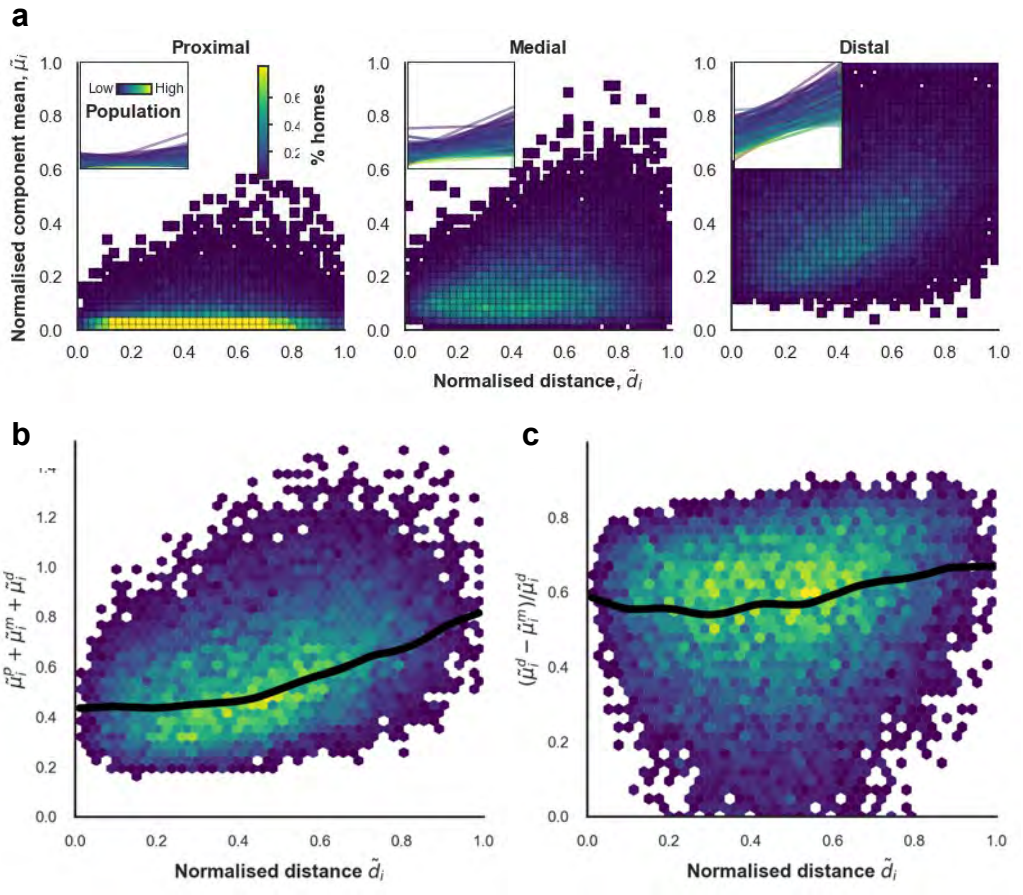

Figure 6.4. **Expanding our analysis.** **a** Normalised relationships pooling 226 European cities; density gradients steepen from proximal to distal and show variation with city size, shown in the insets with linear fits per city, colored by rank size. In these insets, slope declines with city size, indicating that smaller cities have a steeper gradient—and suggesting that smaller cities are more dependent on the urban core. **b** The combined weighted length of proximal, medial and distal components grows with distance to the nucleus: people in the periphery need to move farther than those in the centre. **c** We show the difference between weighted medial and weighted distal means as neighbourhoods get further from the centre, which grows noisily with distance: the distal is growing faster than the medial, which suggests that as commutes lengthen residents find relative efficiency elsewhere.

Discussion

In this brief chapter we introduced a method for modelling human mobility as a mixture of activities according to the distance from home. We decompose each neighbourhood's trip lengths into three Gamma components—proximal, medial, distal—estimated with expectation–maximisation after k-means initialisation on log-distances; the result is a "mobility signature" per neighbourhood: six numbers (means and weights for each scale). We can think of this as dividing activity into that which is anchored around home, work, and an intermediate distance that includes amenity clusters like shopping districts. To quantify core–periphery dependence, we regress component means on distance to the effective centre identified from the busiest hexagon; slopes capture how each scale responds to centrality, while intercepts capture baseline outwardness for central neighbourhoods.

Operationally, with $\tilde{\mu}_i^c \approx a_c + b_c\,\tilde{r}_i$, the intercept a_c is a baseline reach at the centre, while the slope b_c is a "nucleus pull" as origins move away from the centre along radius $\tilde{r}$. This separation lets us see where cities rely on the nucleus, a large b_d, against where

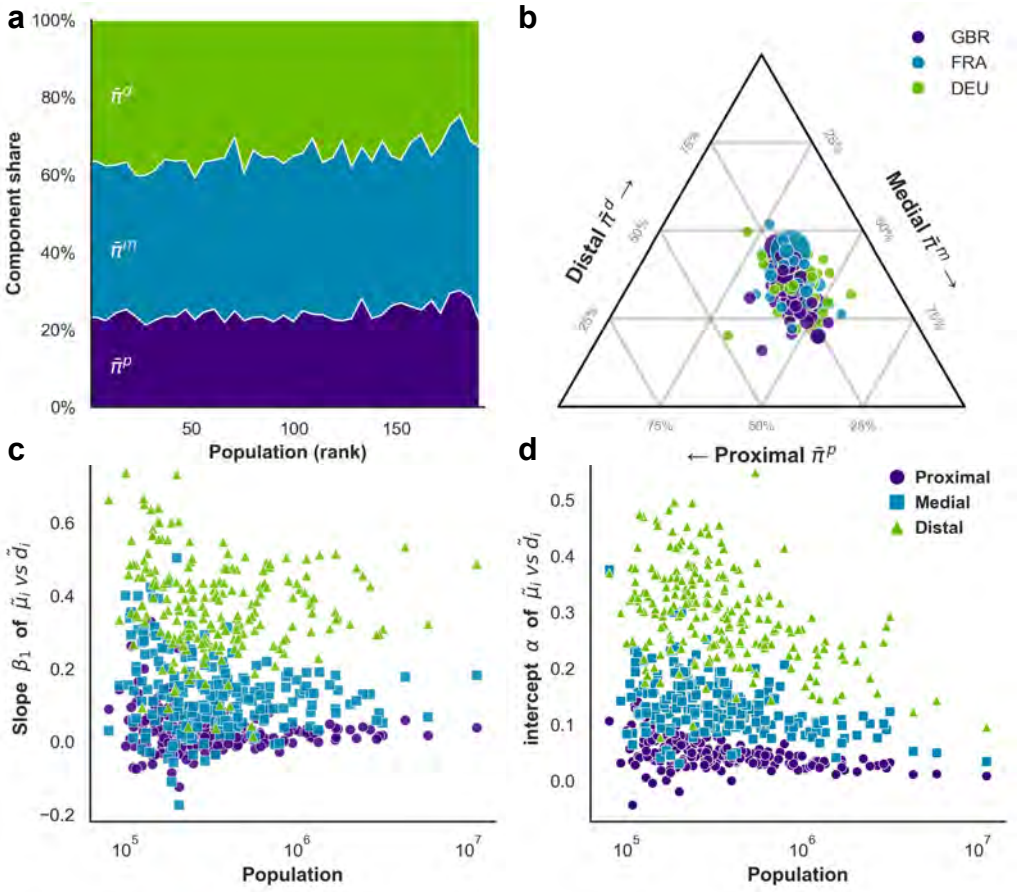

Figure 6.5. **Regularities across cities. a** The weights, or shares, of each component are stable with size: the distal component is typically the largest by weight, comprising $\sim 35\%$ of trips, and the proximal component is the smallest. **b** Ternary composition plot shows no significant differences between countries. Positions encode mixtures where $\bar{\pi}^p + \bar{\pi}^m + \bar{\pi}^d = 1$. Guides are absolute isopleths where for each component. Cities cluster in the interior with an elongation toward the medial axis: distal varies least, while most between—city variation reflects a trade—off between medial and proximal shares. **c** Across cities, the slope of the line regression distal mean against distance is largest but there is no relationship between city size and slope. **d** We do see a negative relationship between the intercept of this regression line and city size, however; this implies that residents of smaller cities traverse a larger fraction of the urban extent relative to residents of large cities.

dispersed opportunities lift activity even at $\tilde{r} = 0$, or a large a_d. Three preliminary findings stand out. First, proximal means display striking invariance with size and location. Second, distal means rise steeply with distance whereas medial means rise more gently, creating a structural distance premium in the periphery. Third, across cities the distal intercept declines with population, while distal slopes are heterogeneous with no robust monotonic relationship to size. The invariance of proximal trips suggests a robust local activity space—an empirical analogue of the 15-minute neighbourhood—whose scale is stable across cities.

Weights and means combine to offer further insights into the structure of a city. Because the mixture always contains a distal tail, peripheral neighbourhoods show inbound distal trips while central neighbourhoods can exhibit outbound ones; incorporating weights into the analysis shows how important each component is to a neighbourhood's budget. Outbound distal components are small by weight but long by length—consistent with occasional long forays to dispersed distal targets within the city. These component-wise gradients are evident in the city maps: central zones show low means across components

and low distal weights, whereas peripheral zones show high distal means. In the pooled view, normalisation exposes the same ordering $b_p \approx 0 < b_m < b_d$, a rising "distance premium" with origin radius, and distal growing faster than medial.

We find that peripheral households pay a systematic distance premium. Let a home cell's distance budget be $D_i = \hat{\pi}_{p,i}\tilde{\mu}_{p,i} + \hat{\pi}_{m,i}\tilde{\mu}_{m,i} + \hat{\pi}_{d,i}\tilde{\mu}_{d,i}$. Because $\tilde{\mu}^m$ and especially $\tilde{\mu}^d$ increase with distance $\tilde{r}$; D_i therefore rises with radius. Cities thus face an efficiency–mixing trade-off: leveraging the nucleus promotes mixing and matching but increases $\hat{\pi}^d\hat{\mu}^d$; strengthening secondary hubs raises $\hat{\pi}^m$ and flattens $\hat{\mu}^m(r)$, reducing peripheral burdens while preserving access to the core. An "efficient polycentric" configuration therefore retains a strong, accessible nucleus for high-value matching while distributing everyday amenities to the medial scale.

Another regularity concerns how intercepts and slopes covary with city size. In our pooled results, distal slopes are generally steepest while proximal slopes are flat, isolating the scales at which cities rely on the nucleus over proximate opportunity. Smaller cities tend to exhibit higher distal intercepts and flatter distal slopes: even central residents undertake distal trips to dispersed targets. Larger cities invert this pattern, with lower intercepts and a wide dispersion of slopes; this implies that in larger cities, residents do not cover as much of the urban extent with distal travel, and contrasts big-city specialisation with small-city generalisation, as residents from all over a small city visit a higher fraction of it.

Across cities, mean distal and proximal shares, $\bar{\pi}^d$ and $\bar{\pi}^p$, vary little, whereas the medial share $\bar{\pi}^m$ does. This motivates a simple "medialisation" index, $\bar{\pi}^m$, that can be tracked over time, compared across cities, or evaluated pre/post interventions—including new rail links, shifts to amenity distributions under remote work. If we are to continue to leverage the urban nucleus for learning, sharing and matching, this medial axis is the primary means by which we can improve urban efficiency.

In previous chapters, we showed that the shift to proximity over centrality might limit socio-economic mixing. Medial locations that bridge between 15-minute neighbourhoods may enable this change, respecting time budgets while encouraging face-to-face encounters. Tracking medialisation over time provides a concrete way to ask whether remote work is producing connected urban villages—or hardening the city into an archipelago of enclaves. With a stable proximal component, changes in the medial component index the strength of the linking nodes between local worlds. With an invariant proximal component, the medial component becomes an important signature of nodes linking otherwise separate 15-minute cities.

This method, expanded from work on transit networks [95] to GPS mobility data, allows for a richer understanding of how people distribute travel budgets throughout a city, conditioning on where they live. Beyond descriptive value, the signature is diagnostic. It yields interpretable indices—e.g., a nucleus pull b_d and an inside–out diffusivity a_d—and an aggregate distance-premium gradient $\mathrm{d}D/\mathrm{d}\tilde{r}$. The method is well suited to equity analyses that identify where D_i or $\mathrm{d}D/\mathrm{d}\tilde{r}$ is highest, and to evaluation of projects that aim to raise $\bar{\pi}^m$ without eroding access to the nucleus. Taken together, the six-number signature repositions monocentricity and polycentricity as points on a compositional spectrum, revealing which levers—nucleus pull (b_d), nucleus diffusivity (a_d), medialisation ($\bar{\pi}^m$), or proximal stability (σ_{μ^p})—actually move realised distance budgets.

Changing Climates, Changing Cities

7 Extreme heat and urban structure

Abstract In this chapter we turn to the threat to mobility from a warming climate, and its possible implications. Extreme heat is a problem globally, with rising temperatures affecting aging populations, but it is of particular relevance in Europe, with lower rates of air conditioning and older populations. Research on mobility during extreme heat remains limited to small samples and isolated contexts, leaving significant gaps in our understanding how entire populations adjust their day-to-day activities and how these adaptations vary across social and economic groups. Here we use open mobility data from passive and active mobile network connections covering 13 million individuals in Spain (27% of the population) to examine extreme heat's impact on mobility at scale. We stratify by age, gender, social and economic status, and activity. Our findings show activity falls by as much as 10% on hot days generally and 20% on hot afternoons specifically, when temperatures peak. Individuals cut infrequent activities most, frequent ones less, and rarely skip work. Further differences emerge on hot days. Older adults are more likely to avoid travel to work and cut back on other activities, while those earning less are less able to avoid work. Our results show that extreme heat disrupts mobility and exacerbates vulnerabilities, affecting older and poorer groups in different ways. These disruptions have significant implications for urban economies, as curbed activity and interaction—both planned and unplanned—may threaten the capacities of cities as hubs of social and economic exchange.

Introduction

Extreme heat poses a serious threat to lives, livelihoods, and the economy [534, 298]. Rising temperatures have been linked to higher numbers of hospital admissions [531, 436] and increased mortality rates [503, 249], with an 85% increase in heat-related mortality for people over 65 between 2000–2004 and 2017–2021 [265]. Extreme heat reduces productivity in both manufacturing [460] and agriculture [445], and slows economic growth [158, 128]. These challenges have been compounded by the increasing intensity and duration of heat waves over the past century [357, 359]; in Europe, summers that were once 1 in 100 year events by temperature may become 1 in 20 year events [290]. Here we develop our understanding of how daily travel responds to extreme heat by linking mobility data, stratified by a rich set of socioeconomic attributes, with climatic conditions. Our findings provide the first estimates of how different populations adapt to heat, revealing constraints and trade-offs while documenting broad disruptions in urban life.

The effects of extreme heat are not distributed evenly across population groups, as some are more vulnerable or exposed than others [534, 253]. For example, the elderly are disproportionately affected by extreme temperatures [521, 347] due to increased vulnerability, often associated with chronic conditions such as diabetes [346, 345], which heighten their risk during heat events [530]. Similarly, while not inherently vulnerable, some workers face significant risks due to prolonged exposure to extreme temperatures, particularly those that engage in sectors of the economy that rely on physical labor [421], such as construction and agriculture [142]. Socioeconomic status is also known to play a key role in determining the impacts of extreme heat, as wealthier households turn on air

conditioning at lower temperatures than poorer households [131]. In relatively affluent regions such as the US, the widespread adoption of air conditioning has significantly reduced heat-related mortality, highlighting the importance of wealth in adapting to climate stress [50, 229].

Individuals respond to hot weather based on their level of vulnerability and exposure to high temperatures, adjusting their behaviour and adapting to the risks posed by extreme heat. For example, time-use studies suggest that individuals reduce outdoor activities and shift towards indoor spaces during heat waves [58, 210], yet such shifts may be constrained by income, occupation, or urban design [210]. Mobility decisions, which we study here, represent a key dimension of behaviour that most people can modify in response to extreme heat [273, 216, 79, 55]. Cities in many contexts see more cyclists and pedestrians on warm days [80], but fewer on hot and humid days [79], demonstrating a parabolic relationship between temperature and active travel. On hot days, more employees report missing work and employees who do work report being less productive [529]. Social and economic incentives may moderate behavioural adaptations, as there is evidence that the risk of heat stroke is heightened in the military and during athletic competitions [356]. Further, shocks to infrastructure may either encourage or discourage adaptation, with railways, roadways, and energy grids experiencing greater strain during heat waves [186]. Such disruptions could, in turn, indirectly influence behaviour, compounding the challenges posed by extreme heat.

While research identifies connections between extreme heat and its impacts on health, behaviour, and economic activity, critical questions remain about the complex ways in which heat shapes human mobility at scale and across different populations [79]. Mobility data derived from digital traces offers a powerful lens to study these behavioural dynamics [224]. Digital traces allow us to monitor how populations respond to extreme weather events, from heat waves to floods [206]. Recent advances in GPS mobility tracking have revealed stark differences in how socioeconomic groups respond to extreme heat [291]. For example, wealthier groups can reduce outdoor activity for longer durations, while disadvantaged groups are often compelled to return to work or other activities sooner, despite continued risks [275].

The European context provides a valuable setting for examining these behavioural dynamics. Unlike many developing regions, much of Europe has historically experienced temperate climates, yet the frequency of extreme heat events has risen sharply in recent decades [526]. This combination of historically moderate temperatures and increasing heat exposure creates a natural laboratory for studying behavioural adaptation. Notably, compared to many other regions, Europeans show vulnerability to heat at lower temperatures [487], making even moderate heat waves consequential for public health and human behaviour. Research shows that around the world, the point at which both cold and hot temperatures cause mortality and morbidity varies with the average temperature [228], suggesting that populations learn to live in different climatic conditions, but it is not clear that adaptation along the necessary cultural and technological dimensions will occur at the speed of climate change in Europe, as recent heat waves have been equally fatal as those in prior decades [48, 196].

Several additional factors make Europe particularly suitable for studying heat-related behavioural changes. First, the continent faces distinct vulnerabilities due to its ageing population, as the elderly are disproportionately affected by extreme temperatures [521, 347]. This demographic challenge is compounded by rising rates of chronic conditions like diabetes [346, 345] that heighten risks from extreme heat [530]. Second, Europe's relatively low adoption of air conditioning compared to regions like the United States [239] means that populations must rely more heavily on behavioural adaptations to cope

with rising temperatures [216]. Third, European cities differ markedly from their global counterparts in their cultural norms around mobility and leisure during periods of extreme heat [16], offering insights into how social factors shape adaptation strategies.

Spain, where we focus our study, exemplifies these European dynamics while facing intensifying challenges. Projections for the next 50 years suggest that Southern Europe, including Spain, will experience a combination of rising temperatures, increased drought frequency, and ageing infrastructure [516]. Heat waves in this region are expected to become not only more intense but also more spatially expansive [281], affecting larger portions of the peninsula during any given event. These challenges highlight the urgent need for behavioural studies that can inform adaptation strategies tailored to European cities and populations.

Here we examine the effect of extreme heat on daily mobility patterns in Spain, thereby shedding light on the behavioural foundations that shape the economic and health impacts of high temperatures. We combine large-scale mobility data collected in 2021 and 2022 with high-resolution estimates of thermal comfort. By focusing on Spain, we explore how extreme heat alters activity patterns in a European context, considering factors such as the type of activity as well as demographic and socioeconomic disparities. We find a clear distinction between routine and non-routine activities: for a given person, work and study patterns change least, visits to places infrequently visited by that person change most, and those to places frequently visited change some but not as much. Activity falls most in the afternoon, when the temperature is hottest, and least in mornings and evenings. Commensurate with greater risk from extreme heat, the oldest populations adjust behaviours more than the youngest, and even travel to work less; further, we identify possible economic constraints, as the less affluent groups change behaviours less than the wealthy. This research offers new insights into how populations manage the challenges of extreme heat, trading off relative risks and needs, and informs adaptation strategies for European cities.

Results

To understand the effect of heat on activity we start by linking data on thermal comfort with data on daily mobility. We then apply several modeling techniques to detect statistically significant changes in mobility behaviour across various population groups.

To analyze mobility patterns, we use data provided by the Spanish Ministry for Transport [314], which contains records for the movements of $\sim$ 13 million individuals, or $\sim 27\%$ of the population. The data represent flows within and between 3,999 districts, including mainland Spain and the Balearic Islands, and are stratified according a variety of important characteristics, including a broad classification of the activity and the trip distance, as well as the age, gender and economic class of the person making the trip. Trips are logged from both active events like texts and calls as well as passive events in the form of probes from the network operator, allowing high temporal and spatial resolution. Because they come from network operators rather than applications, these data have comparably less bias than data from aggregators of GPS location data and they are validated and balanced with surveys and administrative statistics to ensure quality and reliability (see Methods for more detail). Fig. 7.1**A** shows the networks for each month, demonstrating its strong coverage in both urban and rural areas, and we show time series and validation in Figs. C.1 and C.2. While many studies using GPS location data are only able to impute demographic attributes using administrative statistics [321, 153], often aggregated to large areal units that make the estimates crude, the data that we leverage here allow us to decompose travel according to demographic attributes

without resorting to imputation. Because the data are aggregated into districts, we cannot see the precise activity, but each activity is given a category: home, work/study, and frequently- or infrequently-visited places. According to the Ministry for Transport, an activity is "frequent" if a given person visits that place more than once in two weeks, and it is "infrequent" if not.

We measure daily variations in experienced heat using ERA5-HEAT data from the Copernicus program [161], which provides the Universal Thermal Climate Index (UTCI). This metric reflects perceived ambient conditions by incorporating temperature, humidity, solar radiation, and wind into a standardised formula. Focusing specifically on the impact of extreme heat, our analysis is restricted to the summer months, defined as May through September, for the years 2021 and 2022. Fig. 7.1**B** shows the mean UTCI across months, with the strongest temperatures in July and August. In both periods of study, Spain experienced significant heat waves: heat waves occurred in June, July and August of 2022, and while 2021 was milder relative to 2022, and temperatures hit record highs in August.

Heat drives a sharp decline in travel for discretionary activities, with minimal impact on work or study mobility

We use a two-way fixed effects model (TWFE) to estimate the causal effect of temperature on activity by exploiting variation in temperature while controlling for both time-invariant attributes and spatially-uniform shocks through district and date fixed effects. The district controls are necessary to strip out spatial confounds and the date controls are important to adjust for temporal patterns, as Spain sees activity change considerably in August as many people make holidays during this month. Shown in Fig. 7.1**C** (and reported in Table C.1 in the Appendix), infrequent trips fall around 10% and trips to work or school see little change, with frequent trips falling almost 3%. We see no differences comparing trips within districts and trips between districts. Note that Poisson TWFE coefficients are log-semi-elasticities so a coefficient of $0.1 \approx 10\%$, or $100(\mathrm{e}^{0.1} - 1)$.

To identify the form of the relationship between temperature and activity, which will allow us to make predictions across the full spectrum of temperatures, we turn to a generalised additive model (GAM), which fits smooth functions to capture nonlinear relationships in data and in doing so extract effects across different temperatures. In this model, we use province rather than district fixed effects because of the computing demands of fitting a GAM. We fit a cyclic cubic spline by day-of-year to model the seasonality, and include day-of-week and holiday terms to account for variations across days. Combined, these model any variance that is attributable to timing. When we attempt to isolate this continuous relationship between heat and activity in Fig. 7.1**D**, we see that higher temperatures result in lower activity. Although this model disagrees with the TWFE in which activities—frequent or infrequent—are most responsive to heat, which likely comes from using fixed effects at the level of the province rather than the district, the GAM shows that there is a temperature at which activity in both categories peaks—something which will we see further evidence of below. (In Fig. C.3, we show these patterns are consistent when we limit our sample to large cities or small cities, and when we hold out 2022 or 2023.) This suggests that as climate patterns shift over time, some seasons and places will see visits and trips increase as warmer weather generates activity while others will see them decrease as hotter weather destroys activity.

A placebo test looks for confounding variables by shuffling temperature and examining whether our model detects an effect that should not exist. Our data allow us to test both spatial and temporal confounding, and we do so by permuting the temperature either within a district, so that each district on each day is treated with temperatures from

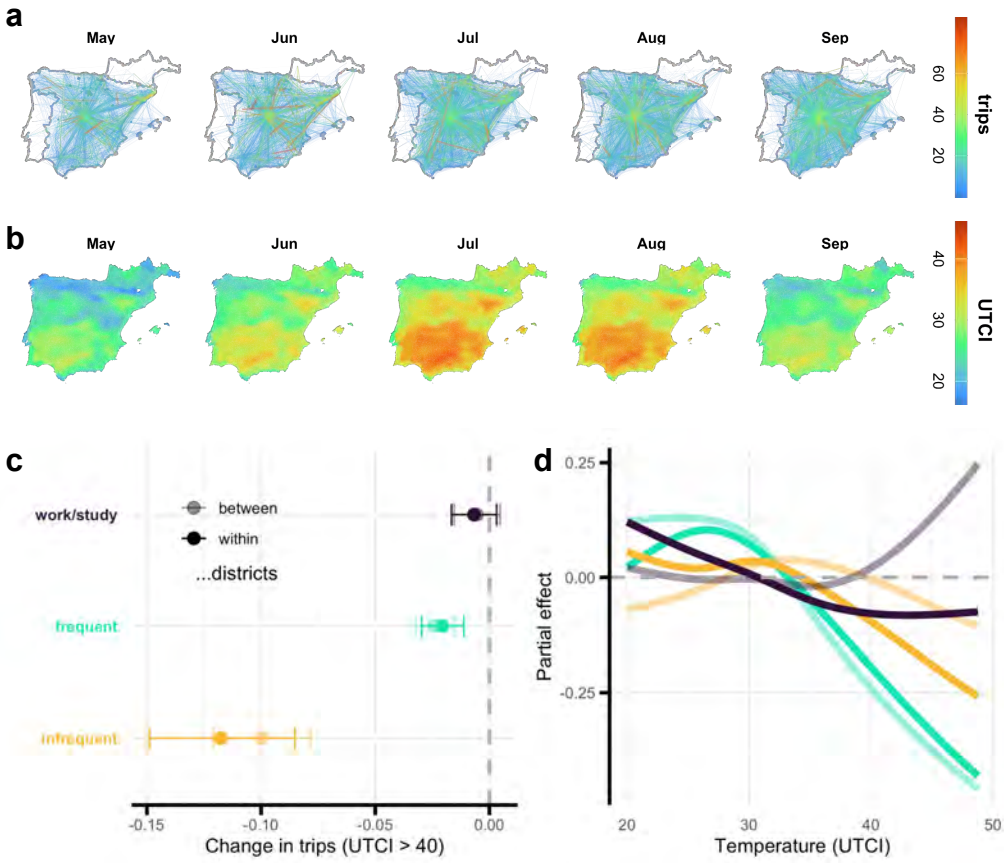

Figure 7.1. **The effect of extreme heat on activity. A** Mobility networks in Spain across months in 2022 and 2023; because the data represent 30% of the population, the networks have good coverage, including rural and urban areas, although the network is dominated by cities like Madrid in the middle of the country. **B** Mean temperatures in the same period, with heat peaking in July and August. **C** Estimates from a two-way fixed effects model, controlling for district and date, on the effect of a UTCI above 40°C, showing that frequent and infrequent activities fall while trips to work or school hold steady. **D** When we examine effect of temperature along the continuum rather than in binary, we see that for both frequent and infrequent activities, and there are warm temperatures that increase activity and hot temperatures that decrease them.

another day of the year, or within a date, so that each district on each day is treated with temperatures from a another district in a different part of Spain. We show these results in Fig. C.4 but we report here that all specifications show no effect of permuted UTCI on activity.

In both models, work and study are not responsive to heat. Further, infrequent trips within a district—those that are more likely to be traversed on foot—fall more than those between districts. Next we leverage the rich demographic and geographic attributes in our data to examine how these effects vary across different populations and contexts.

Extreme heat affects mobility most for the elderly and poor, with no significant gender differences

In addition to activities, we can also disaggregate our data according to various demographic attributes, including age, gender and income. Our preferred specification is the TWFE which, under certain assumptions (see Methods section), allows us to estimate a causal effect of extreme heat on mobility. In the following investigation, we use temperature bins to assess the effects of different temperatures on activity, from mild to hot, rather

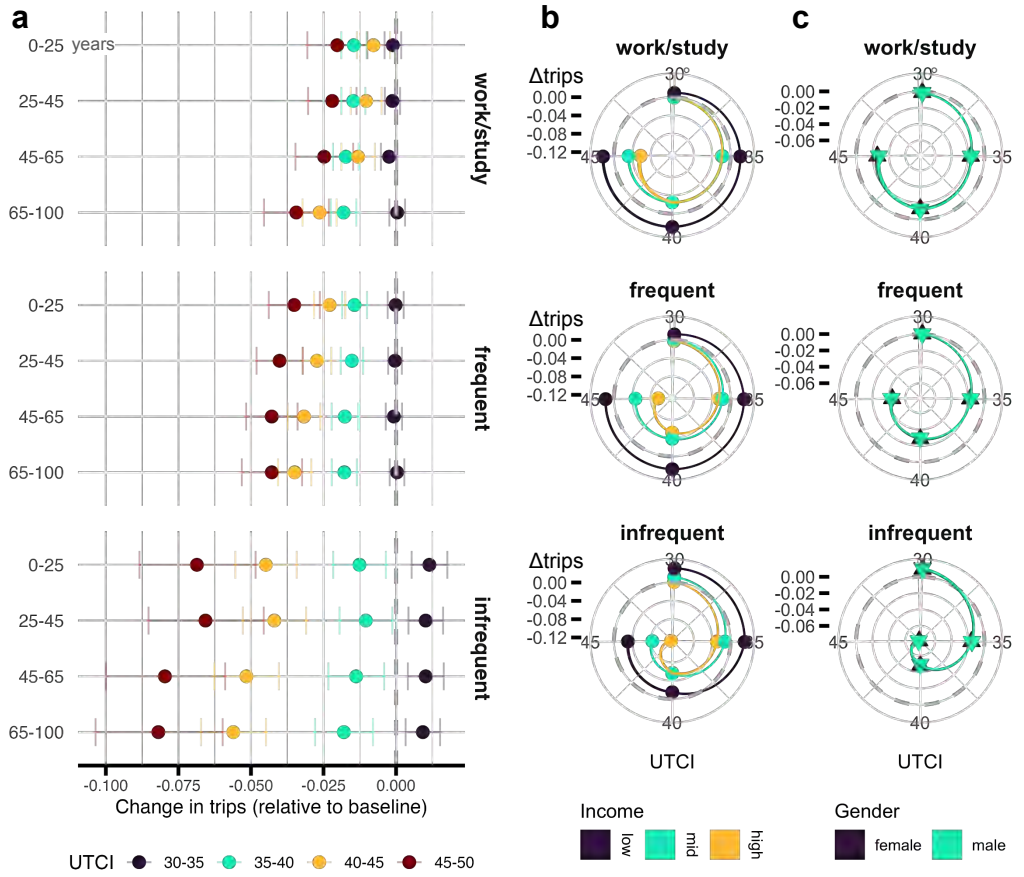

Figure 7.2. **The effect of extreme heat on different groups.** **A** Disaggregating by age, we see that as temperatures increases and the regularity of the activity decreases, trips decline, but more for the oldest than the youngest people. We also note that warm-not-hot temperatures generate infrequent activities, people may expand their repertoire of activities during clement weather. **B** Change in effect size as UTCI increases for different income brackets, showing that the poor less sensitive to high temperatures than the wealthy, which may be tied to work-from-home, although all groups reduce infrequent activities. This suggests that discretionary mobility changes but obligations like work do not. **C** We see no differences by gender, with activity falling identically as UTCI rises for both males and females despite aggregate differences in mobility between the genders.

than a binary indicator of extreme heat. Our estimates are both statistically and practically significant, showing reductions across all classes of activity at various temperatures. Although treatments—different temperatures—are not always statistically distinct, the patterns we see in this section vary systematically between activities and monotonically across temperatures, lending confidence to the relationship we see across point estimates: higher temperatures mean lower mobility.

In Fig. 7.2**A** we see a gradient, with higher temperatures corresponding to stronger declines in activity across all age groups. Looking at how different ages respond to extreme heat, our results are clear: mobility for the young is the least affected by high temperatures and the impact becomes larger as age increases. For the oldest group in our data, a given day with heat index above 45°C corresponds to an 8% decline in infrequent activity, a 4% decline in frequent activity, and even a 3% change in work or study. Because we are using data from 2022 and 2023, we must note that these effects do not necessarily mean a reduction in work, because much if not all of the foregone travel could be to work from home. Yet this change in behaviour would still have implications for cities as Spain warms over time: if the elderly are *missing* work, there is a direct eco-

nomic cost to the workers, but if they are simply working from home, then the economic burden falls on the shops and restaurants that rely on business from commuters.

This is consistent with the fact that heat poses a greater risk to older populations than it does to younger ones [347]. In relative terms, the oldest are most affected by extreme heat but because they constitute a larger and more active population, the greatest decline in absolute terms comes from the middle-aged population, while younger populations are least impacted in both absolute and relative terms.

There are many plausible channels by which income and heat could interact but here we propose two that we then compare to the data: on the one hand, the wealthy might be more capable of coping with extreme heat, via air conditioned homes and cars, and thus remain unaffected; on the other hand, the poor might be less able to afford missing work. There is also a disparity in who can work from home, with many white collar jobs allowing at least some remote work and many blue collar jobs not allowing it at all. Shown in Fig. 7.2**B**, our results accord with the conjecture that the poor cannot afford to, or are not able to, miss work. Another relevant consideration is that work trips here will include leaving and returning after Lunch, so this need not indicate fewer trips to work per se, but fewer excursions from the office. Particularly, individuals from households in the lowest income bracket are unaffected by high temperatures while those from households in the highest income bracket reduce travel across all classes of activity.

Supporting the hypothesis that work compels the least affluent group to stay active, we see that this group still curbs infrequent activities while holding steady trips for work or study as well as for frequent activities even when the heat index surpasses $45°$C. These frequent activities could be attached to daily or weekly routines like lunch breaks, or taking children to school and therefore co-occur with work. For the wealthiest group, all three classes of activity fall at that level of discomfort—by as much as 10% and 15% for frequent and infrequent activities respectively.

With large differences in labor force participation between men and women [340], and differences in both unpaid work and care work between men and women [339], we might expect routines to vary enough to see variation in responses to heat. Instead, once we stratify on the type of activity, we observe no gender differences in mobility. The lack of any difference here is consistent with research that shows relative progress toward gender equality in Spain [84], but because we cannot identify parents in our data, we cannot make any more specific claims. (We also note in Fig. C.5 that there are no statistically significant differences between the trip counts in the network when we stratify on gender.)

Larger drops in the afternoon and on short trips that may involve active travel

Because heat is variable throughout the day, starting off cooler in the morning, heating up in the afternoon and cooling later in the evening, we test whether or not the patterns we see demonstrate that people respond to changes in temperature as the day progresses. Fig. 7.3**A** shows that visits to all classes of activity fall more in the afternoon on hot days than they do in the morning, by as much as 20% for infrequent activities on the hottest days of the year. Yet even frequent activities, which may be coupled with work or study, fall by more than 10%. Taken together, this also lends credibility to our earlier estimates because it shows that mobility responds not just to hot days but the hottest part of the day, which would be less likely if we were observing a spurious effect.

In Fig. 7.3**B**, we explore effects across different journey lengths. Our data do not allow us to interrogate why people might be avoiding certain kinds of travel, but in our models we see that the largest reduction in activity that comes from trips that span less than 2 km, which are less likely to involve a car. This agrees with literature showing that cycling and walking are most impacted by hot days [79]. Again we see tight coupling

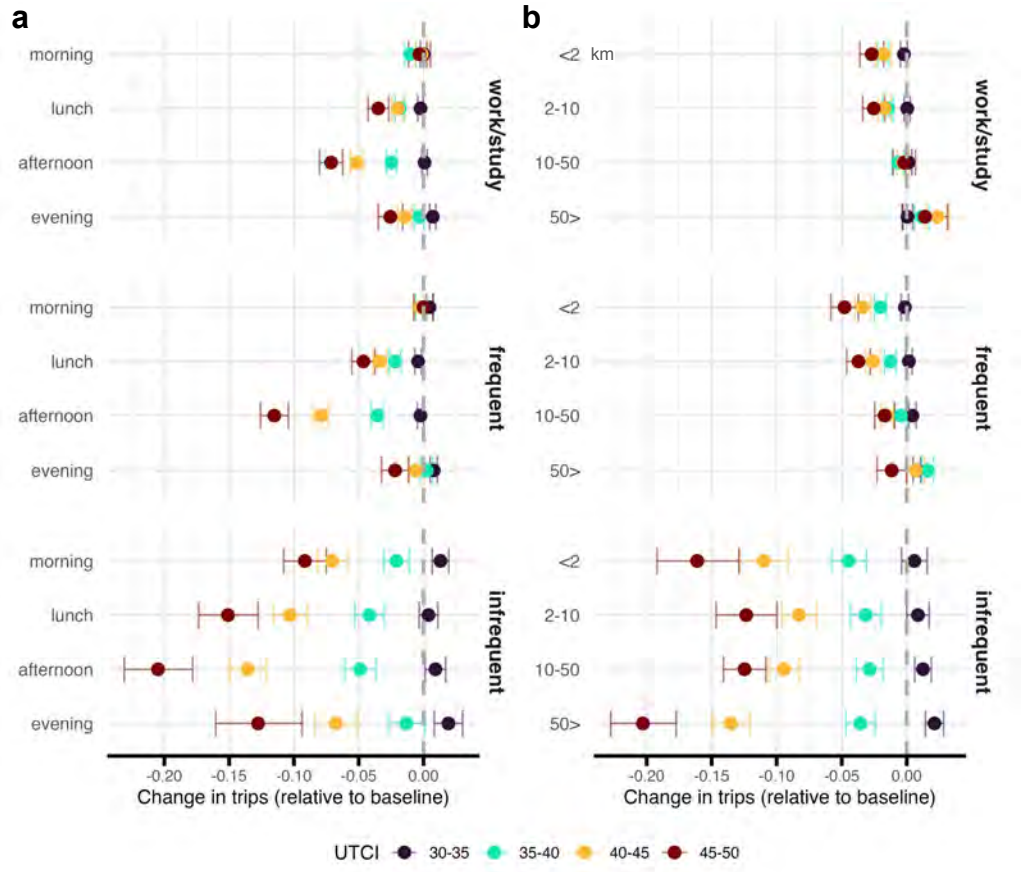

Figure 7.3. **Time and distance.** **A** We see patterns consistent with our earlier findings, wherein trips to frequently-visited locations fall more than trips to work or school, and trips to infrequently-visited locations fall most, but we also see that this effect is pronounced in the afternoon when temperatures are crest. We do not see evidence of substitution to the morning or evening before or after temperatures peak, as those trips still decline on hot days. **B** Generally, longer trips are most resilient to extreme heat and shorter trips, and cars may play a role in this difference as shorter trips are more likely to taken on foot, but infrequent trips see strong changes at both long and short extremes.

between mobility for work or study and for frequent activities, suggesting that certain activities might go hand-in-hand with a routine that includes both professional obligations and personal needs. Although long trips are generally the least affected by high temperatures, they experience the largest declines when they involve infrequent activities, specifically when UTCI exceeds 40°C.

While we again see evidence that long trips for infrequent activities increase when temperatures are mild and decrease when they are hot, we find no evidence that people are moving from hot to cool with long trips. In Appendix Table C.5, we modify a standard "gravity model", which estimates flows between districts based on distance and population, to include the temperature gradient between origin and destination; this gradient shows no effect on flows.

Reduced social mixing as temperature rises

We also document significant changes to the structure of the mobility network as temperature changes, which may have implications for how urban areas function and how social groups mix. Fig. 7.4**A** and 7.4**B** show how trips flow to and from districts with different populations and different incomes (in deciles), respectively. Generally, trips

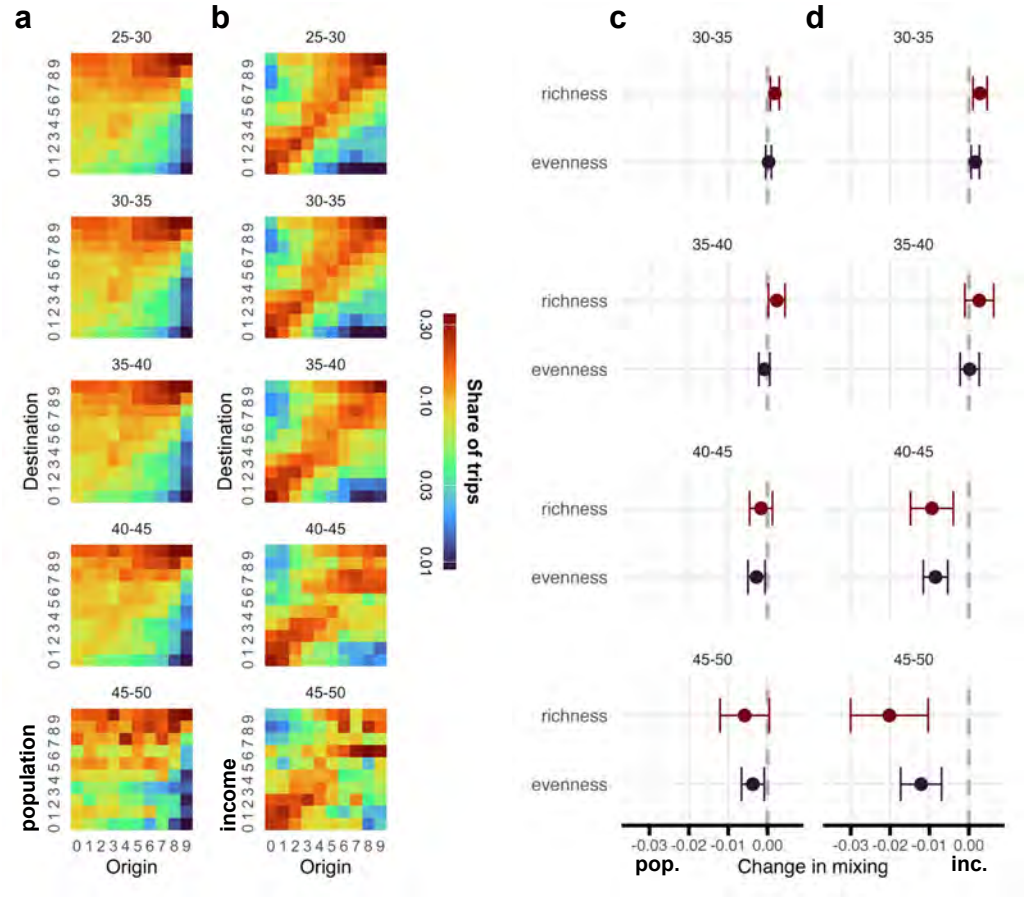

Figure 7.4. **Mobility patterns and heat. A** Changes in flows between districts ordered by population decile, showing that at normal temperatures most flows go from low population to high population areas, or stay within high population districts, but this relationship breaks down during the hottest temperatures and more trips flow between what are suburban densities. **B** Changes in flows between districts ordered by median income decile, showing that there is a subtle but consistent bias wherein people from middle income districts tend to visit upper income districts, and this relationship also fades at higher temperatures—although trips continue to flow between lower income districts. **C** The results of a more controlled test, using a TWFE to predict the change in richness and evenness of visitors to a destination according to origin population, and **D** according to origin income; we see a stronger decline income diversity and the same gradient we see with total activity, with stronger temperatures generating stronger effects.

flow from less populous to more populous areas (urban bias), and from lower income to higher income (wealth bias). These patterns are marked by the higher values in the upper triangles of the matrices in Fig. 7.4**A** and 7.4**B**. We stratify on UTCI to show how these mixing profiles vary under different climatic conditions, and plot results by row according to different UTCI bands. When it becomes hotter, these twin biases attenuate. More trips occur within middle quantiles, and in particular few trips flow from middle income to high income. Agreeing with our earlier results showing limited change amongst the poorest, flows within and between low income districts holds constant. (Note that although districts are large units, with ~ 8 thousand residents, the ratio of *between*-district to *within*-district flows is 3:1, so while many needs are met within each district here is substantial potential for mixing.)

Building on these descriptive results, we introduce another TWFE design to relate mixing with temperature, again controlling for district and date. To do this, we compute two metrics, borrowing from ecology [375]: *richness*, defined as the raw number of income or population groups who visit, and *evenness*, measured using the Shannon

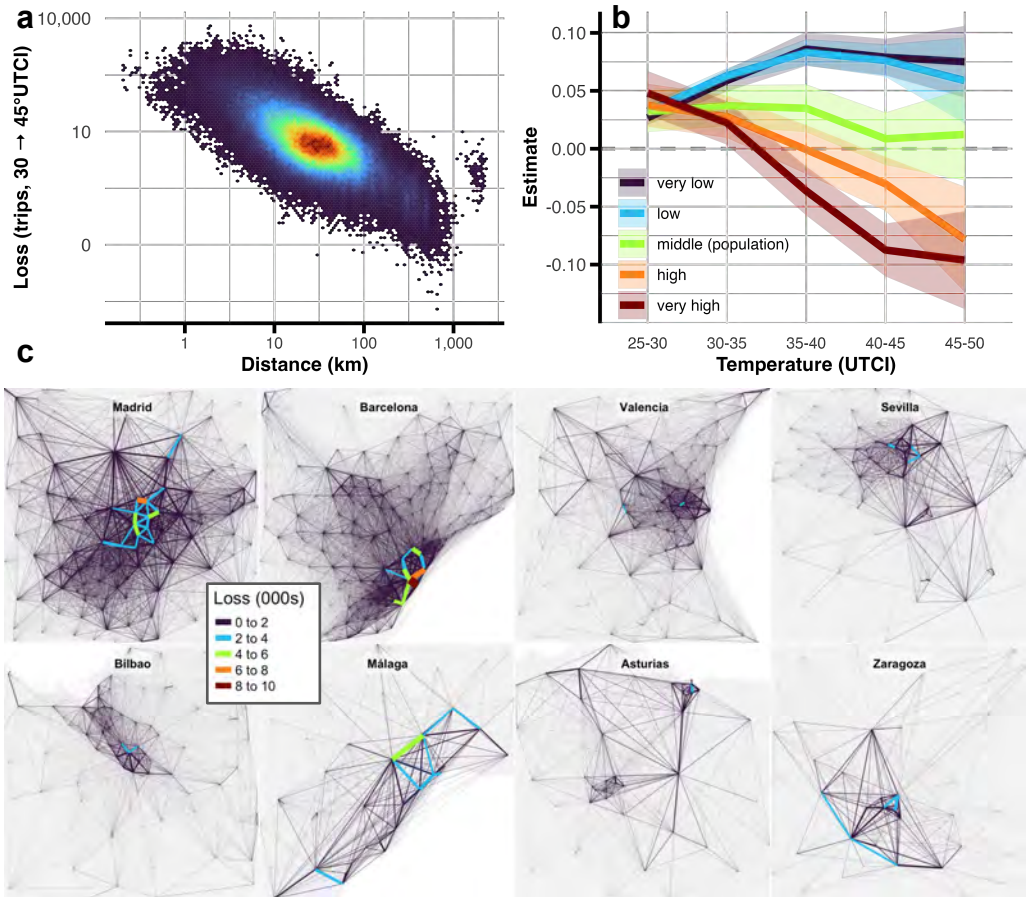

Figure 7.5. **Urban impacts.** **A** We build a network of trips between districts and show the relationship between the loss of trips between districts and the distance between them, according to a gravity model with temperature, finding a clear negative slope. **B** In order to model how different parts the rural-urban gradient respond to heat, we add population quantiles to that gravity model, and see that more populous areas see larger reductions in trips. **C** Building out the networks from this gravity model, we can see that edges in the core see the largest declines in flows, while edges in the periphery are preserved.

entropy of visitor distribution across population or income groups. Note that we are testing for something subtly different, because in the mixing matrices we look at correlation between origin and destination and here we are looking at the diversity of origins present for a given destination. The new metrics, richness and evenness, are used as the dependent variables in our TWFE model. Looking at mixing between low and high population districts in Fig. 7.4**C**, we see less of a change than what is visible in the matrices, although there may be a slight reduction in mixing between rural, exurban, suburban, and urban classes at the highest extremes. Looking at mixing by economic class in Fig. 7.4**D**, however, richness falls by $\sim 2\%$ and evenness falls by $\sim 1.2\%$ on very hot days in this more controlled setting. Taken together, this suggests that much of the change in mobility that we see in the matrices is attributable to seasonal variations, yet it also shows a direct effect from heat on mixing. As we move from lower to higher temperatures, we see a consistent progression; we also see evidence from more mixing on mild days, fitting with earlier indications that some temperatures are conducive to new activities.

Because we see changes to urban structure under this model, we model those changes formally using a gravity model that considers population at origin and destination, along with distance between districts and the temperature on the day. The gravity model, shown in Table C.2 confirms the earlier models but gives us the ability to observe changes on the network. In Fig. 7.5**A** and **B**, we show the components of that model: shorter edges see

flows decline more on hot days that longer edges, and flows between populous areas fall most. Rural areas see mobility taper off from the optimum—at $\sim 35°$C but the net effect relative to cold weather is still positive. We see the consequences of this in Fig. 7.5C, where we see that losses on hot days tend to be concentrated in the urban core rather than the periphery.

Heat steepens effective distance: near links shrink more than far links and populous nodes shed the most flows, hollowing the network's core. Generally, in this network analysis we see that social life in cities will be disrupted if extreme heat worsens without adaptation, with less mixing and less activity in city centres—a development which could threaten the economic and social advantages of cities.

To make sense of our findings, in Box 5 we turn to anecdotes from around the world, but developed and European economies in particular, to see if our work agrees with reports in the media. We see clear evidence that people both *reduce* and *retime* activities. Many countries, including Spain, have adopted workers protections to limit work during the hottest times of day and on the hottest days. This fits our data showing that activity during the middle of the day is most affected. There are also subtler ways that heat could influence activity patterns: cities have started to recommend visits to cooling centres, which could reshape activity and reallocate budgeted time away from retail and leisure on hot days; further, changes to transportation schedules could make travel less convenient, adding a time cost on top of the health risk. Reports discuss "intertemporal substitution"—moving activities to different times of day. This is not wholly inconsistent with our data, which shows declines day-on-day as well as during hot afternoons, because those shifted activities could still come at the expense of others—and days with shifted activities could still have net fewer trips than those without the need to adjust any plans.

Projecting into the future

We use modeled estimates of Earth's future climate, derived from CMIP6 [480], which represents the state-of-the-art in climate projections (see Methods for further detail). The strategy that we employ here is simple: we switch out the UTCI for a given district and day in 2023 for the UTCI on that day in 2073, 50 years later. Holding all else constant allows us to predict activity using our fitted models. The simplicity of this approach introduces limitations, which we discuss now, but it does allow us to understand the implications of our findings in possible future with little adaptation.

Our estimates assume that temperature and nothing else changes going forward. Although the "Lucas Critique" [289] tells us to be wary of making projections when humans can adapt to changing circumstances [245], Spain is ageing and thus the demographic issues that we highlight above could be exacerbated. The exercise is also informative because activity is fundamentally different to mortality and morbidity, where we do see evidence for adaptation in recent years [229, 50]: travel within and between cities will likely require, for the foreseeable future, contact with the ambient air. Coping with heat may reduce mobility more, not less, as we attempt to reduce its worst effects on health. For example, it may become more common for employees to work from home on hot days, as telework has changed the demand for face-to-face interaction [53]. For this reason, we believe our estimates to have important implications for the future economy of cities and towns, including, for example, urban business that depend on office workers or rural areas that depend on tourists [523].

Our model predicts a 3.5% reduction in frequent activities during summer months, May to September, and a 4.7% reduction in infrequent activities. Travel to work or school is projected to fall by 2.2% during that same period. Yet over the full year, all activities will decline by just 1%, as warmer weather during what were once cold periods increases

"Anecdata" on heat and behaviour in advanced economies

Developed economies. Cities and agencies systematically adjust schedules and expand cooling: a bellwether for extreme heat, Phoenix now operates extended-hour and 24/7 cooling sites through summer and is consolidating a permanent overnight hub; libraries and pools extend hours during alerts [124, 333]. Generally, railways and airports reduce operations in heat, which could discourage activity [326, 331, 60]. Employers and regulators adjust work, and new contracts between unions and employers require air conditioning in new delivery trucks [395, 396, 397]; Protections for workers during extreme heat are expanding, and now include both indoor and outdoor working conditions in some regions [338, 100, 36]. Sports and events shift timing/format: for example, tennis matches now activate enhanced heat rules like roof closure, extended breaks [405, 389]; Globally, sports federations formalise cooling breaks and later kick-offs [413]. Utilities and transport agencies issue guidance to slow down during heat spikes as national and local lines post delays [37]. Reporting also documents how delivery/logistics workers alter hours, seek shade, and improvise cooling in vehicles while pressing firms toward greater air conditioning and shift changes [377].

Europe. Governments often suspend mid-day exposure and adjust timing under heat alerts: Greece has closed the Acropolis and shortened outdoor work hours during peak heat [400, 398]; France has shuttered schools and adjusted public services [408]; Italy has periodically limited outdoor work and issued sectoral adaptation measures [410, 391]. British railways have introduced speed restrictions and warn of track damage, with operators publishing dedicated timetables for hot days alongside mitigation protocols [326, 331]. European supply chains reroute, for example, as low Rhine water forces barges to sail at reduced loads and triggers surcharges [392, 412, 411]. Retail footfall can thin in the mid-day heat, with listed chains flagging sales slowdowns attributable to fewer people on the street [337]. A broad synthesis describes the recurring operational pattern: runways and railways buckle, metros close, kitchens shut when staff cannot safely work; firms adopt shift changes, such as shading and cooling from fans or mists [318].

Spain. Spain has moved toward codified work adjustments in heat: national rules enshrine bans on some outdoor work during severe heat alerts; guidance affirms worker rights, including to interrupt activity in the face of heat risk [394, 177]. The health ministry has published a *Plan Nacional* which emphasises prevention/anticipation and targeted communication during alerts [313]. Municipalities use climate shelters and timing shifts: Barcelona has expanded to $\sim$ 400 *refugios climáticos* with coverage targets by barrio [176, 11, 174]; Madrid publicises discounted private venues and identifies hundreds of public spaces as shelters, while sometimes restricting park access under extreme conditions [175, 178]. Transport operators activate, which include providing access to drinking water on trains, during heat waves [388]. Behavioural reporting shows mid-day suppression and evening substitution: Iberian households shift errands to later hours; beaches, malls, and shaded public spaces absorb demand when temperatures run 10°C above seasonal norms [401]. Jobs in construction and manual labour adopt summer schedules or mid-day pauses [215, 104].

Box 5

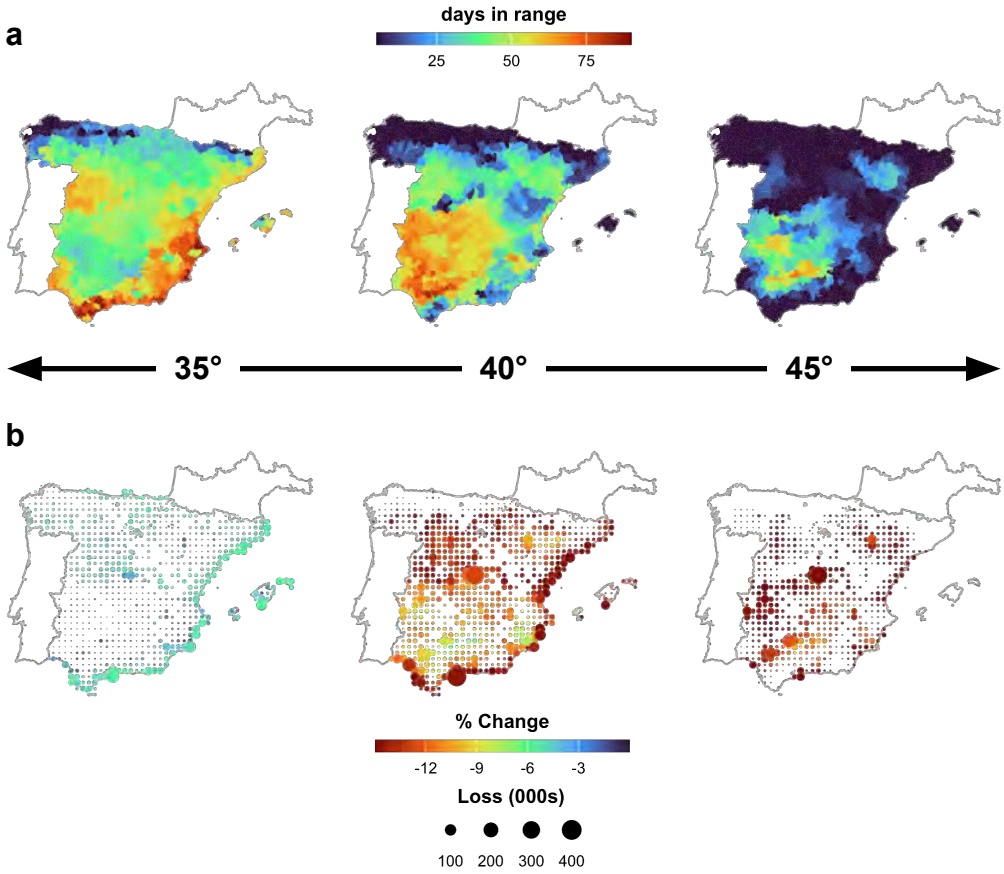

Figure 7.6. **Projecting in the future. A** Prevalence of temperature exceeding a given value. Maps show that Spain's South will see 75 days between $40°$C and $45°$C as well as dozens of days above $45°$C; days between $35°$C and $40°$C will also be more common in the North. **B** Consequences of those temperatures in both relative and absolute terms per district, showing that according to our model the largest effects at lower temperatures will be in areas that do not experience high temperatures now. Cities in the South like Sevilla and Malaga will experienced the highest temperatures, however, they do not see the strongest drop in relative terms; big cities like Madrid and Barcelona will see the largest change in absolute terms.

mobility while hotter weather during already hot seasons decreases it.

Next we explore how this could vary across cities and regions. Fig. 7.6**A** shows the number of days spent in a given UTCI range, from $35-40°$C to $40-45°$C to $45-50°$C. All of these temperatures correspond to reductions in activity according to our models, and they will be common in the South of the country. In particular, parts of Andalusia will see temperatures exceed $40°$C for more than 75 days each year. The North of Spain will experience more days that exceed $35°$C, but will only have a few days a year at the extremes that the South will endure.

Fig. 7.6**B** shows that the effects of these changes may change mobility if adaptations do not address rising temperatures. Our estimates here show the change on a given day at a given temperature, compared to the baseline. Effect sizes are larger at lower temperatures in the North, where society is not as well acclimated to heat. Because extremes that will occur in the South will not be prevalent there, we concentrate on temperatures between $35°$C and $40°$C; days in this range would still reduce activity by as much as 6% on days when they occur. While the South is largely unaffected by these temperatures, it see strong changes to activity when temperatures exceed $40°$C; when this occurs, activity in the South will fall by 6-12% while activity in the remainder of the country could fall be 15%.

Discussion

Extreme heat deforms the urban social reactor: mixing profiles flatten and core–core edges lose the most volume, implying fewer cross-class and cross-district encounters precisely where density normally amplifies interaction. The diversity of visits falls, both with core-periphery dynamics and with socio-economic mixing, indicating less heterophily. Short trips—most likely walked—contract first in the afternoon. This could combine to erode the incidental exchanges and spontaneous encounters that underpin both retail activity and weak-tie formation.

Because older and wealthier groups curb discretionary mobility most on hot days, lunch-time demand in job-dense cores thins even when work/study trips remain comparatively stable. In short, these patterns have implications with the mechanism that we have discussed in this thesis so far, for cities as social reactors and places of exuberant diversity: the conditions that generate superlinear spillovers—dense, diverse, repeated contact—are precisely those most sensitive to mid-day heat [242, 70, 68].

We find that mobility behaviours respond to heat in ways that are consistent with expectations rooted in the literature on extreme heat and health. First, daily movement patterns for the elderly are most affected by extreme heat; second, afternoons see the greatest decline in mobility levels across population groups, as temperature—and thus risk—crests. Commensurate with heightened risk, we find that those older than 65 are more likely to reduce activity in response to high temperatures, and that their reductions are greater at higher temperatures. Compatible with the daily temperature dynamics, we see that on hot days activity falls most during the early and late afternoon, and least in the mornings and evenings. Again, our results show that as temperature increases, so does the effect size in our model estimations.

Yet we document opposing effects from extreme heat amongst another vulnerable group. While the oldest dramatically reduce discretionary activities and skip travel to work, possibly to either work from home or miss work, the poorest do not miss work and reduce activities less. This means that the group most at risk, because age magnifies the threat of heat, is responding according to that risk, but it also suggests that the poor are least able to compensate for extreme heat by foregoing work and travel. This reveals important economic constraints that may influence mortality and morbidity.

Lending confidence to our estimates and findings here is the consistency with which our models behave: higher temperatures correspond with larger effects in all of our specifications. Although these findings appear intuitive, our study is the first to document these changes accounting for trip and individual attributes; in doing so, it demonstrates the adaptive nature of human mobility in the presence of extreme heat: populations respond to high temperatures by changing routines and avoiding certain activities. Without more granular data, we cannot shed light on what all of these activities are specifically, but we do see patterns in the analysis we are able to make.

With an ageing population and a warming world, our results suggest that policies to adapt to extreme heat will be important for keeping Spain and possibly other European countries active and productive in the coming decades. Yet many existing strategies to mitigate the worst health effects of extreme heat involve air conditioning [50], which is difficult to apply to activity. For example, cooling shelters can relieve much of the threat of heat for vulnerable populations [502]. To be clear, there are still strategies to mitigate extreme heat between buildings, like greening [244], and certain modes of travel can also be air conditioned, but a broad drive to move between air conditioned spaces could change the social fabric in cities, and we find preliminary evidence of how this might occur here. The short trips that are more likely to be walked are also more likely to

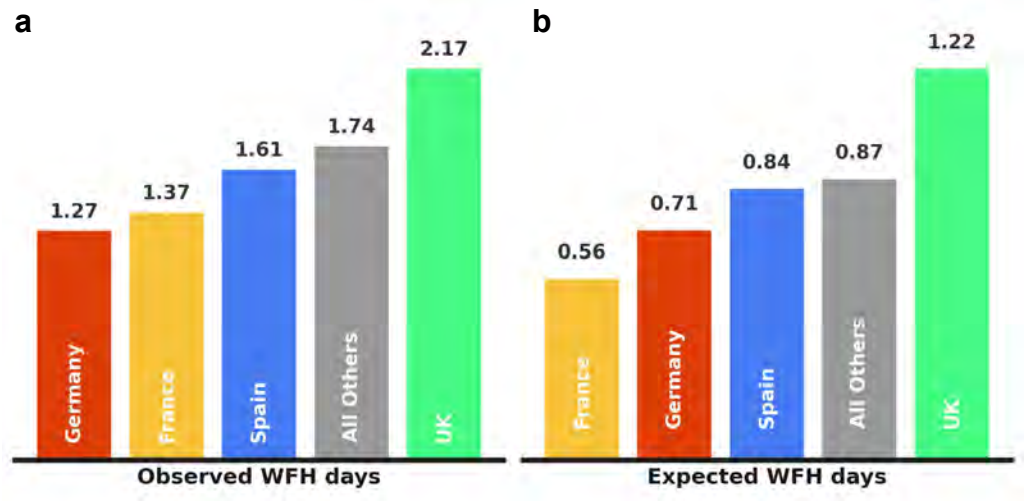

Figure 7.7. **Working from home.** Data from Wave 2 of the Global Survey of Working Arrangements [12] indicate that remote work is in line with the European average in Spain. **A** Among full-time, prime-age, college-educated workers—the group most likely to be able to adopt work from home—Spanish respondents reported an average of 1.61 days worked from home per week. **B** Employer guidance points to a contraction over time, with expected work-from-home days falling to 0.84. In contrast, employers in the United Kingdom report both higher observed levels and less severe expected reductions, suggesting that Spanish institutions are comparably less accommodating of remote arrangements. The level of remote work matters for adaptation to climate stress: if Spanish workers must commute even during heat, the burden of adjustment shifts from individual behavior to collective and infrastructural responses, reinforcing the need for institutional adaptation.

be avoided in extreme heat. On the hottest days, many also avoid travel to work, so businesses that depend on commuters for foot traffic might suffer. In particular, remote work may become an alternative to commuting on hot days. This suggests that the very adaptations that make heat survivable might erode the subtle interactions that make cities engines of innovation [32, 41] and culture [202]. Policy, therefore, is not only about cooling individuals but about preserving contact in cores (hours, shade, cooled access) to maintain the agglomeration benefits that heat otherwise erodes.

Data collected in 2023 [12] suggest that remote work is similar in Spain relative to other European countries. Using a sample of full-time, prime-age, college-educated workers to survey those jobs most likely to adopt work from home, these data [13], shown in Fig. 7.7**a**, indicate that workers spend 1.61 days per week worked from home. Guidance from employers, however, suggest that this number will fall over time: Fig. 7.7**b** shows that expected days worked from home in future will fall to 0.84. This could put an institutional constraint on the rational choice to avoid hot commutes, and while it will create discomfort it will also preserve the social reactor. Further, if workers *must* go to work, it puts pressure on cities to adopt further adaptions—instituting collective approaches instead of relying on individual responses. Although our results are consistent, intuitive and suggestive, the mechanisms underlying the behavioural changes we see in the data during extreme heat are not clear. For example, we cannot prove a link between short trips and active travel, and although we know that frequent trips appear linked to work/study, which agrees with research on trip chaining [316], we cannot determine which classes of amenities are most affected when people avoid travel to work. Our study thus limited, providing strong evidence for adaptation while much of the details are left to speculation.

Finally, our work demonstrates the value of open mobility data: the Spanish government makes the mobility patterns used in this study freely available and further software development has made it easily manipulable [260]. Despite aggregation and suppression

to preserve anonymity, our results are clear and robust. More governments should make data of this nature available in order to allow researchers to advance our understanding of human mobility in a changing world. While the data here reveals important insights into who adapts to heat, as well as how and when, it is still limited by how activities are classified and locations are aggregated. With access to data on human mobility expanding over time, future work should replicate this work in other contexts and decompose mobility patterns according to more specific activities, occupations and locations, which will allow governments and philanthropies to target cooling efforts at those most exposed and least able to compensate for extreme heat.

In the following chapter, we will see similar mid-day suppression but different incidence: poorer areas in India and Mexico contract most, while Indonesia shows smaller elasticities—different deformations of the same underlying mechanism, as heat raises the generalised cost of moving around a city.

Methods

Data

We use daily origin-destination data provided by the Spanish Ministry for Transport [314], which integrates anonymised mobile phone records with demographic, land use, and transport network information to produce a mobility data product. This dataset captures trips over 500 meters within Spain and infers key travel characteristics, including origin and destination points, travel modes, and trip purposes—either work/study, frequent or infrequent locations for that device. The Ministry leverages state-of-the-art algorithms to transform raw mobile network data into structured and scalable matrices, offering high-resolution insights into mobility patterns across spatial and temporal scales.

The data are stratified to allow for interrogation variations across demographic groups and trip purposes [449]. Activities are classified based on the Ministry of Transport's recurrent mobility analysis, which tracks origin-destination pairs over 2-week periods. Destinations visited more than once in this window are classified as frequent activities, while those visited only once are considered sporadic activities. Balance is achieved using official statistics from the national statistics agency to account for differences in age, income, and regional population distributions. This ensures that the dataset is representative of the broader population, minimizing biases associated with the uneven distribution of mobile phone users. The integration of demographic and geographic information also allows for the segmentation of mobility patterns by municipality, province, and other spatial units, providing a flexible foundation for granular analysis.

The Ministry for Transport employs rigorous quality controls to ensure the reliability of these data [472]. Anomalies in travel patterns are monitored through automated systems, which compare data against historical trends and predefined thresholds. Possible errors, such as geolocation inaccuracies or missing records, are flagged and investigated to ensure data integrity. The data are also validated with independent sources, such as FAMILITUR survey data, to confirm the consistency of observed trends with government statistics. Additionally, logical consistency checks, such as evaluating the symmetry of origin-destination flows, are conducted to ensure that the data align with expected behaviours. These efforts, combined with transparent methodological documentation, make this dataset an important resource for understanding mobility in Spain.

We link these data on mobility with an index of thermal comfort from ERA5-HEAT climate reanalysis data [161], provided by the Copernicus program. Universal Thermal Climate Index (UTCI) combines temperature, wind, radiation and humidity to measure

not just how hot it is but how it feels—for example, if humidity limits evaporative cooling through perspiration. We use zonal statistics to compute the mean UTCI at 16:00 for each district on each day.

Models

We employ twin modeling strategies to understand the relationship between heat and mobility, the first to measure the causal effect and the second to estimate the functional form. Both assume the number of trips T terminating in district i at time t follow a Poisson distribution such that $T_{it} \sim \text{Poisson}(\mu_{it})$. Our first approach uses a two-way fixed effects model (TWFE):

$$\log(\mu_{it}) = \beta(\text{UTCI}_{it} \times \text{activity}) + \alpha_i + \gamma_t$$

where μ_{it} represents the expected number of trips, α_i represents district fixed effects controlling for characteristics of the *place*, while γ_t captures date fixed effects accounting for patterns common across districts at a given *time*. The interaction with activity type allows us to estimate differential temperature responses across activities. This specification leverages within-district variation in temperature after accounting for common temporal shocks, providing causal estimates under the assumption that temperature variation is as-good-as-random after controlling for location and time fixed effects. We cluster standard errors at the province level to account for spatial correlation in the error terms.

Our TWFE model enables us to make causal inferences if certain conditions are met, but there are threats to the model. A key assumption is that mobility patterns would have evolved similarly across districts in the absence of temperature shocks (parallel trends). This is plausible in our setting because temperature variation is as-good-as-random after conditioning on location and time fixed effects, and districts cannot select into heat "treatment". A more serious threat to identification would be if extreme heat causes people to substitute between districts, violating the stable unit treatment value assumption (SUTVA). We test for such spillovers using a gravity model of bilateral flows between districts, explained and shown in Table C.2 in the Appendix. After controlling for origin and destination populations and distance, we find no evidence that temperature gradients between districts drive mobility patterns, suggesting SUTVA violations are unlikely to bias our estimates.

To explore and model the potential curvilinear relationship between heat and activity, we complement the TWFE analysis with a Generalised Additive Model (GAM):

$$\log(\mu_{it}) = f(\text{UTCI}_{it}) \times \text{activity}_i + \beta_1 \text{popularity}_i + \beta_2 \text{province}_i + f(\text{DoY}_t) + \text{DoW}_t + \text{holiday}_t$$

where $f(\cdot)$ represents a cubic regression spline with 4 knots, $f(\cdot)$ is a cyclic cubic smooth for the day-of-year, capturing seasonality and drifts in the data, and we control for mean visitation (popularity) and geographic variation (province). We also add day-of-week and holiday fixed effects because, for example, weekends and holidays might have different levels of activity and this allows the intercept to vary on those days. While the TWFE isolates the causal effect, the GAM reveals the functional form of behavioural responses to temperature variation through its flexible smooth functions. The GAM's strength lies in its ability to detect and convey nonlinear relationships without imposing *a priori* assumptions about the functional form, allowing us to identify potential threshold effects and complex response patterns in human mobility.

Projections

In order to understand possible futures according to our models, we use climate scenarios derived from state-of-the-art “general circulation" models [480], which are used by the Intergovernmental Panel on Climate Change to make projections about Earth's future climate. (For context, Fig. C.6 shows projections for major cities.) We extract data on temperature, humidity, radiation, and wind to compute UTCI manually for the year 2050 under the $2 - 4.5°$C of warming, using the simulation from the Centre National de Recherches Météorologiques in France, because it shows high accuracy against observed data [139] and of comparably accurate models it performed best in Spain when we checked it against our data for 2022 and 2023.

Our strategy for estimating mobility in the future is simple yet crude: we replace the UTCI for a given day-of-year in 2023 with the UTCI for that same day 50 years later in 2073, predicting with the new temperatures and all else equal. We compare predicted values using 2023 temperatures to predicted values using 2073 data to ensure that we are comparing like with like, modeled estimates in both cases, rather than using observed values in one and predicted values in another. Because the GAM allows us to produce estimates across the full range of temperatures, we make our predictions using this model rather than the TWFE.

8 Climate, heat and adaptation

Abstract In this chapter, we look at the effects of heat on populations around the globe to its broader consequences. As discussed, extreme heat is a growing threat to both individual livelihoods and broader economies, killing a growing number of people each year as temperatures rise in many parts of the world and limiting productivity. Many studies document the link between heat waves and mortality or morbidity, and others explore the economic consequences of them, but few are able to determine how populations respond to the shock of extreme heat in day-to-day activity. Toward this end, we investigate the link between human mobility and ambient temperature. Examining Indonesia, India and Mexico, we show that extreme heat reduces mobility by up to 10% in urban settings, with losses concentrated mid-day. We examine the shape of the relationship, finding that while heat reduces activity, very hot days and very long heat waves may induce more of it, indicating different adaptation. Effects are stronger in poorer areas. Twinning these models with climate projections, we show that without adaptation mobility may fall 1-2% per year on aggregate, with certain seasons and places seeing activity fall by as much as 10%. According to our estimates, small cities will face the highest relative losses and large cities will experience the greatest absolute impacts.

Introduction

As we noted earlier, extreme heat threatens lives globally [534, 298], with rising temperatures linked to greater mortality [249] and morbidity [531]. Here we take a global focus on behavioural change, with an emphasis on time allocation across hours and short trips—because these are the channels through which heat enters the urban production function and attenuates realised agglomeration [210, 93]. Beyond human health, heat waves induce economic damage in developing contexts by lowering productivity [460] and slowing growth [158]. The severity of this challenge continues to mount: heat waves are intensifying in both magnitude and frequency [357], with their duration growing since 1950 and that rate of growth accelerating since 1980 [359], and much of this change will be felt in the countries we turn to now—India, Indonesia, and Mexico. Yet, the behavioral dimensions of heat impacts are not clear in these and other developing countries, and research is often limited by context, often focusing on advanced economies [350]. By pairing weather and socio-economic data with GPS mobility, we show that extreme heat fundamentally disrupts activity urban areas. In doing so, we demonstrate a scalable approach for modeling human adaptation to extreme heat across diverse contexts.

The consequences of extreme heat differ based on who people are and where they live. Age, gender, and surrounding conditions shape vulnerability [521, 347], while geographical features like water bodies [90], urban design [236], and broader geographic characteristics [537] alter heat resilience. Socioeconomic status also matters: in areas with abundant air conditioning or car ownership, heat waves pose fewer disruptions [216], whereas disadvantaged communities experience sharper mobility shifts due to hot weather [216]. In the US, wealth reduces heat-related mortality via air conditioning [50], as people pay more to avoid heat than cold [16], cut working hours, and redirect outdoor leisure indoors

[210]. Moreover, although wealthier populations can better delay activities, poorer ones resume hazardous routines sooner, facing heightened exposure [275].

Populations do adapt over time: mortality decreases with more frequent high temperatures, implying compensatory mechanisms [516], and the least deadly temperatures often match local climate norms [526]. Governments now classify extreme heat alongside hurricanes and cyclones, going so far as to name heat waves [310] and devise "heat action plans" [228]. Yet adverse outcomes persist: extreme heat lowers productivity in manufacturing [460] and agriculture [445], dampens growth in poorer regions [158], and a single anomalously hot day can curb output [532]. Heat also contributes to conflict [91], erodes learning in non–air-conditioned classrooms [351], and strains infrastructural [186] as well as natural systems [476]. Consequently, understanding how people adjust their daily activities is crucial.

Heat exerts its effects via physiological stress—causing dehydration, fatigue, and organ damage, especially among older individuals [305, 228]—and by compromising cognitive performance [146]. This leads to fewer working hours [529], constrained productivity in physically demanding sectors such as construction [421], and altered urban behavior. Rising temperatures reshape activity patterns in cities worldwide [273, 291], driving individuals to forgo short trips [216], shift destination choices [234], and reduce both public transport use [334] and walking [263]. These behavioral modifications, rooted in discomfort or safety concerns, underlie broader economic and social impacts. As climate change intensifies, deeper knowledge of how heat influences collective movement has become a research priority [501].

In this chapter, we investigate how extreme heat modifies urban mobility in three developing countries—India, Indonesia, and Mexico—home to 1.8 billion people and spanning eight climate zones [63]. We combine mass GPS-based mobility data with thermal comfort measures to quantify the form and magnitude of heat's influence in 2019. By modeling changes in trips, distances, and destinations, we illuminate the behavioral foundations of adverse health and economic outcomes. Given that hotter regions suffer the greatest productivity losses under climate change [142], understanding these responses is essential. We find that extreme temperatures significantly alter mobility patterns, with activity levels dropping by up to 10% during heat waves. We further disaggregate our analysis by population density and socio-economic status, yielding comparative insights into how different groups manage and recover from extreme heat across diverse climates and cultures, showing that economically disadvantaged areas show stronger reductions in mobility. While populations attempt to compensate by shifting activities to cooler hours, this adaptation is less effective in poorer areas, likely due to limited access to cooling infrastructure and inflexible work schedules. Our findings reveal how climate change could exacerbate existing inequalities in developing countries, where the most vulnerable populations face the greatest disruption to their daily activities. Since urban economies depend on the mobility and interaction [171], our results suggest that without interventions to help populations adapt to extreme heat, rising temperatures will impose a cost that falls most on those who can least afford it.

Results

Impact of hot days on mobility

We use daily flows between origin and destination geohash5 aggregations, comprising ~ 236 million trips and ~ 1.4 billion unique GPS signals, gathered from mobile devices in the course of daily usage (for descriptive statistics and illustrations of the network,

see Fig. **C.7** and **C.8** in the Appendix). As these data best represent urban populations, we remove rural areas—clipping each network to "functional urban areas" [441]; Because the network is sparse and our concern is "activity", we define it as the number of journeys ending in each geohash5. For each geohash5, we acquire Universal Thermal Climate Index (UTCI) [161] from ERA5-HEAT [133], which takes wind, sun, moisture, and temperature into account to produce a measure of how hot or cold it feels. Joining the datasets, we are able to make inferences about the effect of temperature on moves terminating in a given geohash5. (See Methods for greater detail.)

We show temperature fluctuations across India, Mexico and Indonesia throughout the year in **8.1A**, finding that each country gives us insights into different climates: located on the equator, Indonesia has a stable climate; Mexico and India have more pronounced curves with distinct profiles. Mexico City suffered heat wave in July 2019 (3+ days in the highest temperature decile [358]), and we show in Fig. **8.1B** mean activity per grid cell during the heat wave compared to the weeks before and average, for comparison. Although noisy, activity falls in the center, with mixed effects in the periphery. With preliminary evidence that heat changes behavior, we use a TWFE approach to understand the effects across all countries systematically. Because frequent temperatures are less lethal [516, 526], we examine deviation from the mean temperature in a geohash5. The results are shown in **8.1C**, which plots the change relative to the strength of temperature anomaly. The results are clear in Mexico: higher temperatures mean fewer trips in a cell. In India, heat reduces afternoon activity—when temperatures peak—but activity shifts to the evening. Sensibly, the effect is inverted for cold, where waiting for the air to warm is better than during extreme heat. In Indonesia, there is limited variation and the effects we can observe are *de minimis*, save for the possibility—given large standard errors—that cold appears to decrease activity throughout the day while heat appears to increase it.

Insights into the shape of the temperature-mobility relationship

To capture incremental temperature changes, we model the temperature–activity relationship using a generalized additive model (GAM), which offers non-parametric flexibility and parametric interpretability. Our GAM uses smooth functions that accommodate varying coefficients and suits the documented nonlinear temperature effects on productivity [92]. (See Methods.) We use a directed acyclic graph (DAG) to identify the minimally sufficient adjustment set—the controls that allow us to infer causal links between our independent variable of interest (temperature) with our dependent variable (activity). For controls, we identify temporal and spatial factors as determinants of activity and use geohash5 fixed effects along with splines to capture trends over time. In alternative specifications, we interact temperature with population and deprivation to explore how these moderate its effects. While the GAM captures the shape of relationship, we use another TWFE specification with various temperature bins to provide more rigorous identification by accounting for both location-specific characteristics and time-varying factors.

We compare both models in Fig. **8.2**, finding consistent results. In Fig. **8.2A** we see that in India, activity declines during the afternoon but recovers in the evenings. Notably, afternoon temperature peaks have inverse effects: most active on cold days, when warmth is beneficial, but least active on hot days. In Mexico, shown in **8.2B**, higher temperatures are associated with fewer trips during all periods of the day. We see limited effects in Indonesia in Fig. **8.2C**, with warmer temperatures increasing activity and cooler temperatures decreasing activity. The highest temperatures have point estimates below zero, but we cannot reject the null hypothesis here.

Mobility displays a concave response around a local thermal optimum, with mid-day and short trips most elastic, consistent with movement penalties and exposure avoidance.

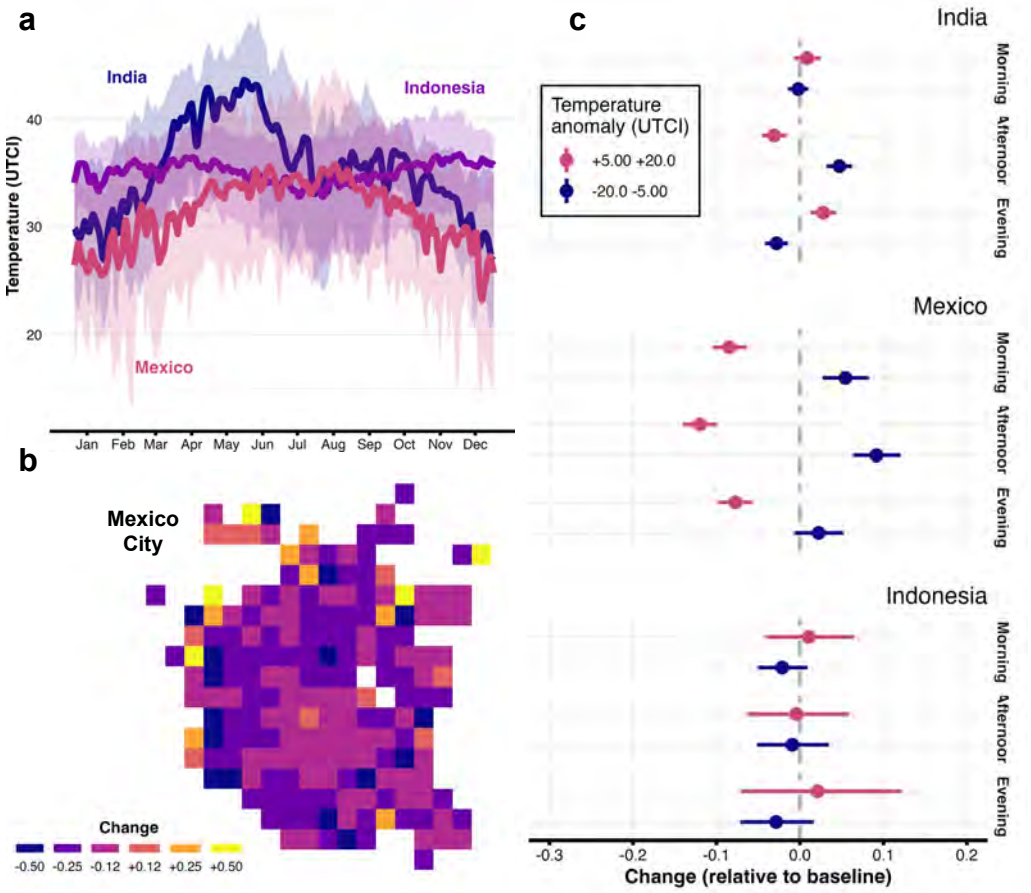

Figure 8.1. **Extreme heat and mobility. A** Temperatures in India, Mexico and Indonesia throughout 2019, showing that we have three different climate profiles: equatorial—and thus stable—in Indonesia, and tropical but varied and seasonal in India and Mexico. **B** Change in mean daily visitation in Mexico City between a heat wave in July and the weeks preceding and following it; activity levels fell 10% on aggregate, and although some peripheral areas gained foot traffic, losses are concentrated in the central city. **C** Results from a TWFE regression showing that, with controls for location and date, activity falls most in the afternoon on hot days—when the temperature is hottest—in Mexico and India. In Mexico activity falls throughout the day, while in India activity rises in the evening, after it has cooled off, on hotter days. Activity rises in the afternoon on cold days, suggesting that on cold days, the warmest hours of the day are conducive to activity, while on hot days, the hottest hours are destructive.

Generally, deviation from a location's mean temperature—hot or cold—reduces activity. Showing that relative changes in temperature are better predictors than absolute changes, our results are consistent with the literature that suggests climate preferences are set locally [16]. This also agrees with research that finds that the most frequent temperature in an area is often the least fatal for that area's population [526].

The public health literature has identified a "minimum mortality temperature" [526], which is the temperature at which all-cause mortality is at its lowest in a given place. In an indication that populations adapt to local conditions, this temperature is typically the "most common" or modal temperature in a place. To explore the possibility of a "maximum activity temperature", the point at which temperature as most conducive to activity, scale and centre (z-score) the temperature in each GeoHash5 so that we consider deviations from the average. Fig. **C.9B** and **C** show the relationship between changes in temperature and trip counts and durations (measured in minutes) for full days. Our results agree with this existing literature from public health: as temperature increases relative to the local average, the number of trips falls. (Although, because the maxima

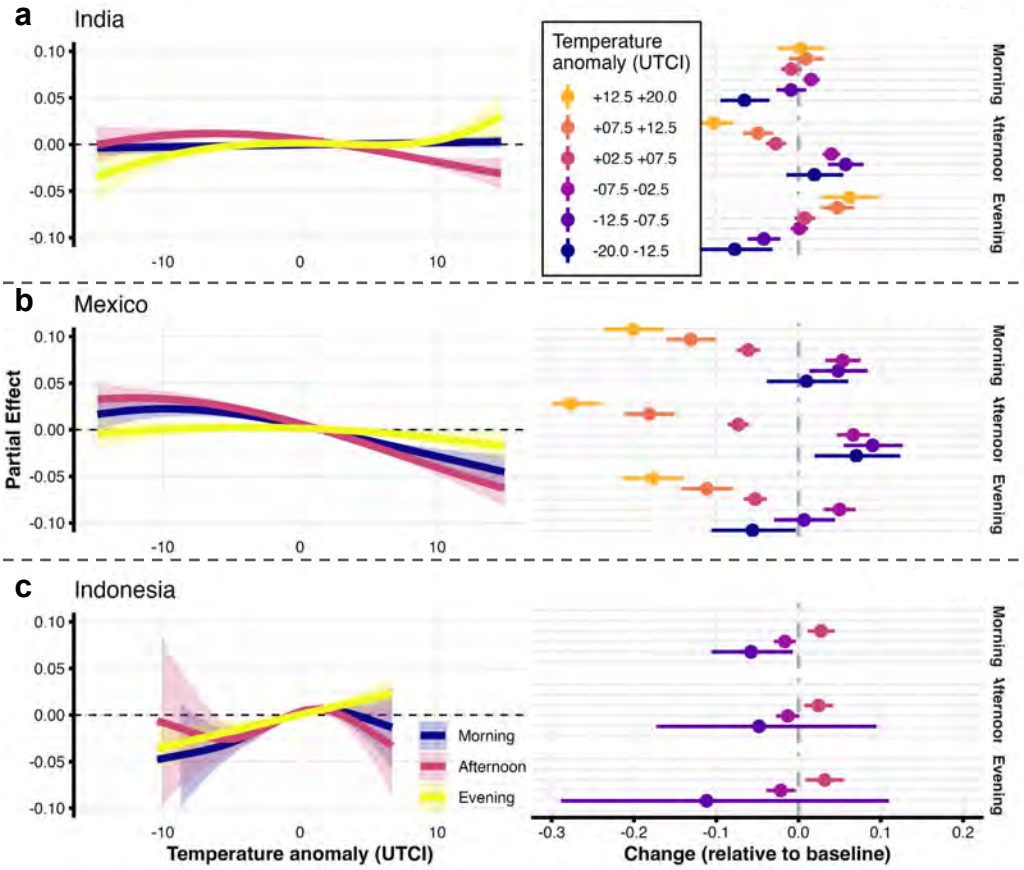

Figure 8.2. **Modeling temperature and activity. A** We decompose the effect, using both a GAM and a more robust TWFE, by time of day, and see that activity actually recovers in the evening in India. **B** In Mexico, a comparably cooler place that might have fewer adaptation strategies, activity does not recover in the evening. **C** The effect is mixed in Indonesia, as an equatorial country with little variation in temperature and possibly strong adaptation strategies; although we cannot reject the null hypothesis, the effect of heat on activity does turn negative at high temperatures. Note also that in all countries, mild temperatures appear to be conducive to activity.

are not at 0 precisely may be due to the fact that we only use the average for one year of weather data.) We also conduct placebo tests in **C.10** and sensitivity analysis in **C.11** to show that our results are unlikely to be spurious.

We then look at absolute temperatures to understand if levels are also a factor in driving activity. To see if prolonged heat waves differ, we test rolling averages of different lengths, where higher averages indicate stronger or longer fluctuations. Fig. **8.4A** shows that in India, when the temperature exceeds $35°$UTCI, the effect of temperature on activity becomes negative; Fig. **8.4B** shows that in Mexico, this change from positive to negative occurs at $30°$UTCI. In India, persistent heat waves (rolling averages above $45°$UTCI) show relative adaptation, with effect sizes rising above zero. While not conclusive, this suggests that extended periods of high temperatures necessitate a return-to-activity, or that policy and behavior adapt only after multiple days of extreme heat. Replicating what we saw in the TWFE analysis, When we partition the data by time of day in Fig. **C.9**, we see that the effect converging toward 0 at $42°$UTCI in India at high temperatures is the product of countervailing forces: extreme heat increases activity in the evenings while decreasing it in afternoons, when the temperature is highest. This difference is only present in India, where temperature is more variable, but it indicates that compensatory changes are at play at the extremes. In Mexico, all times of day see

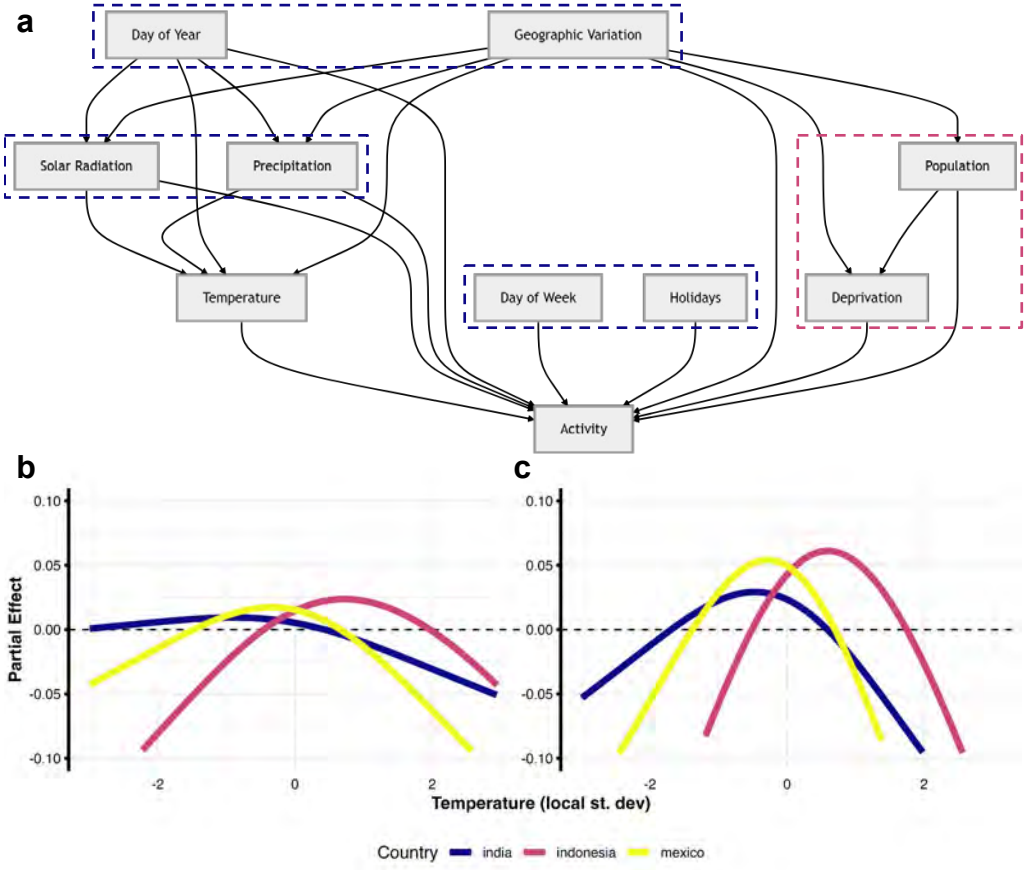

Figure 8.3. **Modeling temperature and activity. A** We set up a directed acyclic graph (DAG) to ensure that we close all necessary causal paths to activity; we control for day-of-year, day-of-week, holidays, solar radiation, precipitation, and use geographic fixed effects that necessarily stratify by population and deprivation. **B** Model results for different countries, showing that the highest activity levels occur at average temperatures for an area, and that high extremes correspond with fewer trips. **C** Model results using trip duration rather than trip count reveal that the longest trips tend to occur at average temperatures, with extreme high temperatures leading to shorter trips.

reductions at high temperatures.

Heterogeneities in the impact of heat on mobility

We next examine how demographic and geographic factors shape activity. In this specification, we use geohash5 fixed effects but add interaction terms for deprivation [191] or population [440]. Building a model that interacts temperature and deprivation, we find that in both India and Mexico, hotter days suppress activity more in high-deprivation areas. In particular, we explore the space of model predictions across all combinations of temperature and deprivation. These predictions are shown in Fig. **8.4C** and **D**, for India and Mexico, respectively. We use the same strategy for population in C.12.

In India, we recover the same arc as above, where activity partially recovers at the highest extremes. The contours in **8.4C** are flatter for low deprivation and steeper for high deprivation, indicating that much of the worst effects of heat in India are felt by the poorest. Looking at Mexico in **8.4D**, higher temperatures bring lower activity for all levels of deprivation, but the contours at higher levels of deprivation indicate that the effect is stronger for poor areas than for wealthy ones. We see recovery at extreme heat in India but not in Mexico, where contours steepen at higher temperatures. This has

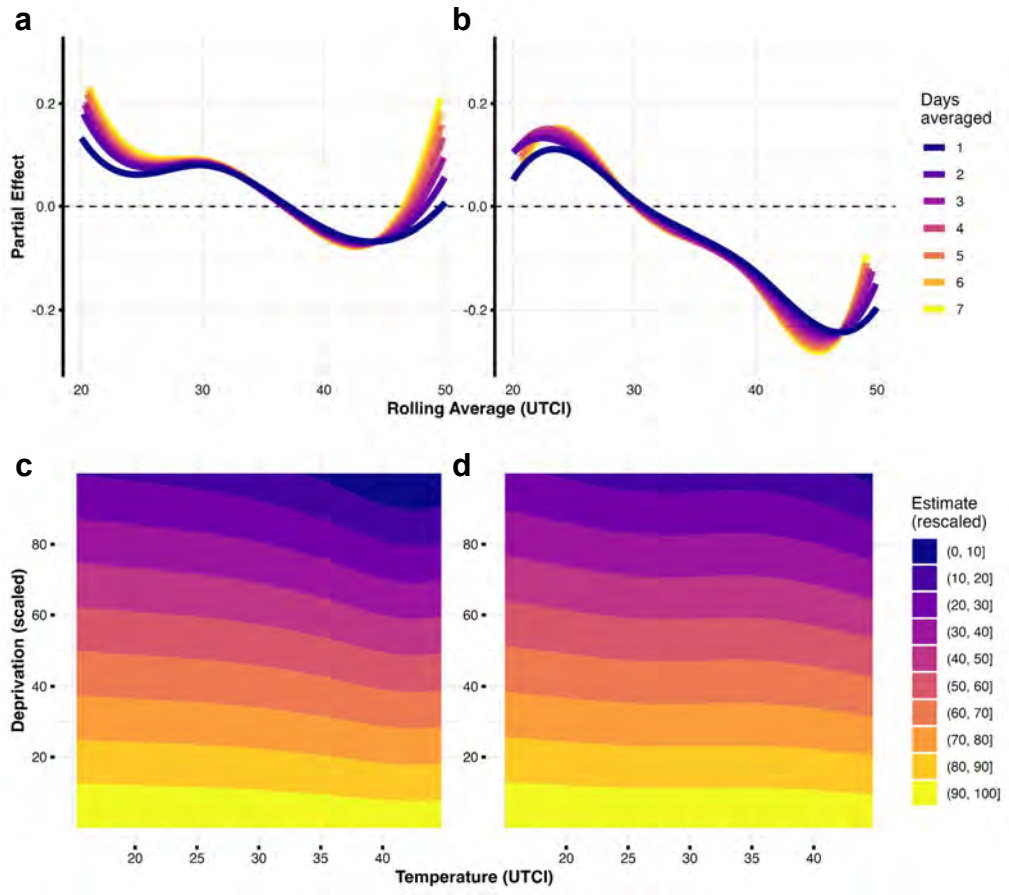

Figure 8.4. **The effect of extreme heat on activity. A** Focusing on India, given that country's propensity for heat waves, we measure the duration and intensity of extreme heat by calculating a rolling average, finding that higher temperatures for longer are associated with stronger effects between $40 - 45°$UTCI but not above that, which may indicate compensatory behavior at a certain point. **B** We fit the same model for Mexico and find similar results, with an ideal temperature and declining activity when heat surpasses it (noting that Mexico does not have the same extremes). Lasting heat waves, measured in rolling averages, show a stronger impact in Mexico and the rebound we see in India never becomes positive. **C** and **D** Looking at the interaction between deprivation and temperature in India and Mexico, respectively, we see a curve that follows above modeling for all socio-economic strata but it is steepest, indicating the strongest effects, for the poorest strata.

important implications for vulnerable populations, as we might expect competing forces to be at play: the poor are also less able to afford taking time off from work, but they may be less able to access cooling, which could suppress activity. We see evidence that the latter pressure dominates the former, and activity is suppressed at extremes.

The sparsity of the data preclude us from making stronger inferences about what is happening to the structure of the mobility network, as we did in the previous chapter. Yet we can use these models to gain important information that feed into structure. For example, because we can see in Fig. **8.4** that the poorest areas see the steepest declines at the highest temperatures, we can infer that the topology is not stable—some groups, and thus some flows, are more sensitive. Likewise in Fig. **C.12**, because the most populous areas have the most to gain on mild days and the most to lose on hot days in absolute terms, again we see that the core-periphery structure of cities is threatened by extreme heat.

As a robustness check and to explore spatial variation, we use an ARIMAX model at the geohash3 level (to ensure full time series data for each). This model predicts ac-

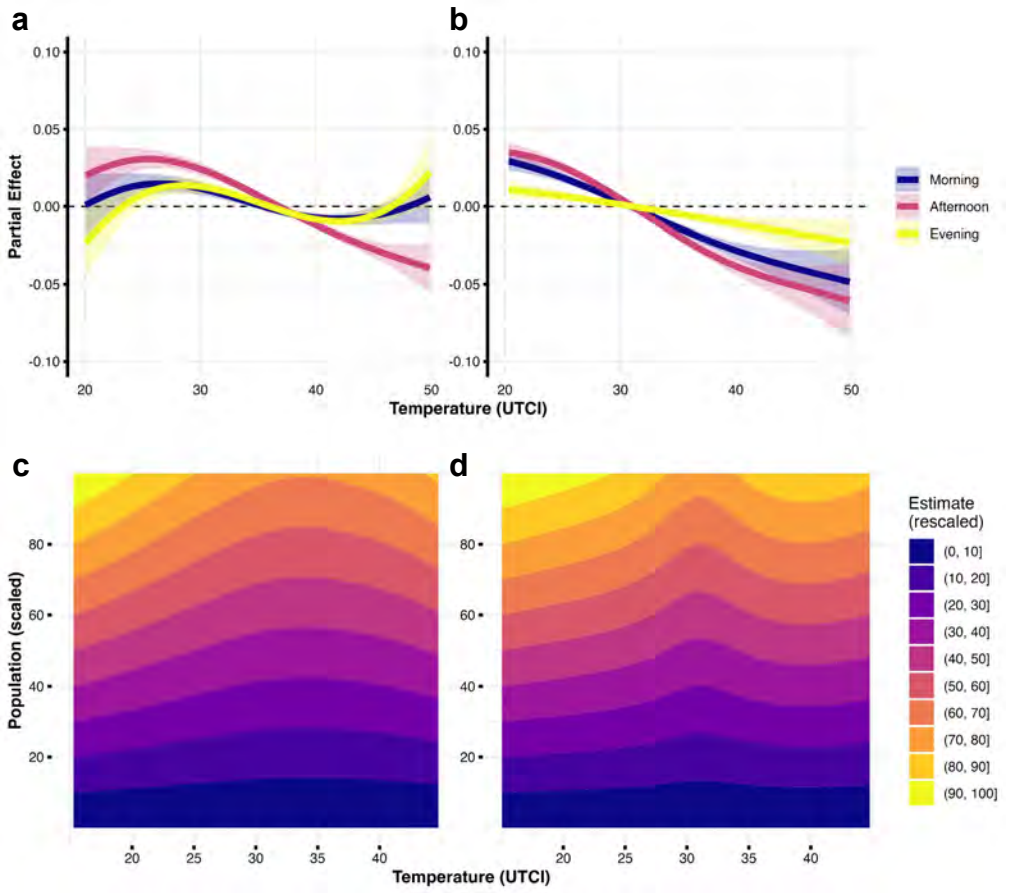

Figure 8.5. **Mobility by time of day and population. A** When we decompose the effect of temperature on activity by time of day, we see that the reduction in activity comes during the afternoons, suggesting that people are avoiding the hottest part of the day. **B** In Mexico, all periods of the day see reductions in activity. We see that in India **C** and Mexico **D**, larger populations see stronger temperature effects in absolute terms, with higher predicted values and moderate temperatures and lower predicted values at the extremes, while these curves are flattened for smaller populations.

tivity by considering activity at time t using activity at $t-1$ and $t-7$ along with linear and nonlinear terms for temperature, which allows the relationship to curve as temperature increases. This method implicitly controls for seasonality and day-of-week effects by using the previous day and the same day of the previous week to make a prediction, and achieves a good fit, shown for India in Fig. **8.6A**. Important to this study, the quadratic term is negative across most regions, indicating that at a certain point the effect of temperature on activity turns negative. Fig. **8.6B** shows, again for India, the sign and magnitude for each geohash3, which shows mixed results but less of temperature effect in the South, where temperatures are more consistent, and more of an effect in the North and Northwest. Some coefficients are positive with high magnitude, but these tend to be in the mountainous North, which does not have as many hot days, and the tropical South, which has less variability in temperature. We take the average coefficient for Temperature^2 across all countries in Fig. **8.6C**. Note that we need to rescale the trip distribution in each country for the effect sizes to be comparable, so the value of each coefficient is difficult to interpret (in Fig. **C.13**, we plot the curves from each ARIMAX formula to show how Temperature^2 influences the prediction). We filter out geohashes according to the length of the time series, because very short series could bias the results; generally, as this threshold of minimum observations rises, so too does the effect size.

Finally, we identify the optimal temperature by making predictions across the spec-

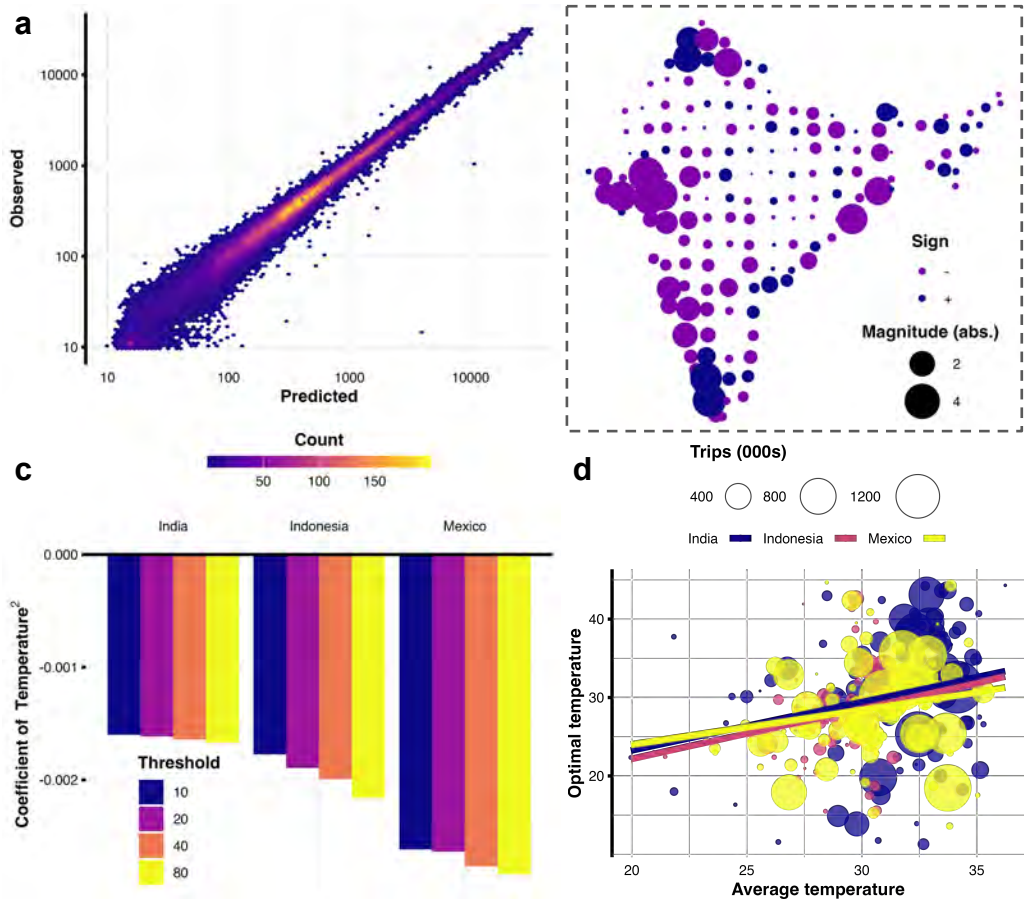

Figure 8.6. **Alternative estimations across all countries. A** We fit an ARIMAX model to the time series of each geohash3, achieving good fit with a mean absolute percent error of 20%. This model uses both autoregression—recent experience—and exogenous variation, temperature in our case, to predict activity, allowing us to interpret the effect of temperature; **B** shows the effect of Temperature^2 across India, with most areas showing a negative effect. **C** As a sensitivity analysis, we vary the required length of the time series, because some areas do not have complete data, and find that the effects only become stronger as we increase the minimum number of days in the series. Although a weak correlation of $\rho = 0.25$, in **D** we see a relationship between average temperature and the optimal temperature for activity according to our model, supporting our earlier results.

trum for each geohash3 and observing where the predictions begin to fall. According to these predictions, the mean and median optimal temperatures for India are 31.1°UTCI and 31.8°UTCI—in line with our estimates from the GAM. (These mirror what we see in **C.13**, which plots the prediction curves for each geohash3 ARIMAX.) Although not a perfect match with the above findings, these results strengthen our assertion that extreme heat can limit activity. We document a noisy relationship between temperature optima across geohashes and the mean temperature of those geohashes ($\rho = 0.25$ across all countries, $p < 0.05$) in Fig. **8.6D**: as mean temperature increases, the optimal temperature for activity also increases. This provides further evidence that populations adapt to climatic conditions.

To make sense of these results, we gather anecdotes and reports from around the world, 40 news articles in total, in Box 6. The picture is clear: government polices often discourage activity during hot days and individual choices respond to changing incentives. The combination appears to be autocatalytic: reduced work and errands—either by choice or by mandate—both causes a direct shift in activity, but it also reduces the incentives around opening stalls and participating in the dampened economy; further,

both costs and risks rise, as workers must move to cooler modes of transit, adding a fee to a journey would be free if walked, and as heat raises the prospect of health complications—especially for vulnerable populations.

Projecting into the future

As in the previous chapter, we next turn to models of Earth's climate [480] to make estimates of how activity will be impacted in future. In particular we use the same data from CMIP6 simulations, extracting the estimates from the Max Planck Institute because these estimates show the best fit to past trends in the regions we study here [139]. We substitute each 2019 temperature-day value with its 2050 counterpart from these climate models. We make these projections knowing that our own study suggests that optimal temperatures are often related to the climate of the area, and thus that the population of a given area may be able to adapt to local temperatures. Nevertheless rapid deviation from these temperatures as global warming accelerates [359] might still be a shock to communities who have acclimated to past temperatures. Furthermore, our data suggest that poor areas suffer the worst effects, and these areas might be the least able to cope with changes to the local climate going forward. As a modeling exercise, this provides us with a worst case, but a plausible one at that.

CMIP6 projections (**8.7A**) indicate 100–150 annual days above $40°$UTCI in Northern states (e.g., Rajasthan), and (**8.7B**) suggests 25 days above $45°$UTCI. We distinguish between large urban areas (top 10 FUAs by population) and small cities (remaining FUAs). On any given day in June when such temperatures are most likely, not a weekend nor a holiday, our model estimates that a single day at $40°$UTCI temperatures to reduce activity in large cities areas by 2% and in small ones by 6%, shown in Fig. **8.7C**. Because megacities generate more activity, they experience a small relative change but a large absolute loss. As we see in Fig. **8.7D**, these losses grow to 4% and 8% for large and small urban areas, respectively, during a single day of temperatures at $45°$UTCI.

We explore the same projections in Mexico, which show the same strong effects from extreme heat and mild effects from temperatures that will occur often. An important difference to note is that while the Southern, tropical areas in Mexico will have the highest average temperature, with Fig.**8.8**A showing most days over $35°$UTCI by 2050, but the desert near the Northern border—where much of the country's manufacturing happens—will experience more heat waves, highlighted in Fig. **8.8**B. Because the literature indicates some adaptation to heat, these heat waves in the North could disrupt activity more than persistent high temperatures in the South; although heat waves will be less common in temperate regions, these areas are most at risk of disruption when heat waves do occur.

In Appendix Fig. **C.14**, we show similar results for Indonesia, with many days over $35°$UTCI in Sumatra and Borneo in **A**. While temperatures of $40°$UTCI will be rarer, shown in **B** they are forecast to happen despite never appearing in our 2019 data. Because Indonesia is in the tropics and does not experience much temperature variability, our model projects strong effects for heat that exists beyond the current distribution. Indonesia is planning to move its capital to the island of Borneo, from Java, by 2045. While this megaproject intends to curb the effects of subsidence, temperatures in this new region, according to CMIP6 data, will experience more hot days than in the existing capital region. Shown in Fig. **C.14C** and **D**, the consequences of these temperatures, according to our modeling, are similar to the effects we see in India: small cities areas with larger relative effects and larger urban areas with larger aggregate disruptions.

A key caveat in year-round modeling: warming formerly cold periods can boost activity, so in India substituting 2050 weather yields only a 1% drop overall. Because

"Anecdata" on extreme heat and behavior

Global. Across regions, authorities and firms shift activity away from the hottest hours and into cooler windows: tourist sites shorten or suspend mid-day opening; governments in chronically hot countries impose or strengthen mid-day outdoor-work restrictions; and employers formalize breaks or remote work during heat alerts [399, 367, 409, 394, 406, 494, 393, 371]. Major events and transport operators *retime* or *reduce* operations, from cooling breaks at the sporting events to speed restrictions and delays on rail lines [407, 390, 326]. Many cities expand access to cooled spaces, while gig applications and delivery firms adjust operations or incentives—all consistent with mid-day labor supply falling and trips shifting to mornings/evenings [369, 370, 484]. Wealthy countries show a wider range of adaptive strategies than poorer countries, for example providing cooled buses to gig delivery workers [483].

India. People substitute into cooled, grade-separated transit and scale back mid-day foot travel: in 2024, the Delhi Metro recorded all-time-high passenger journeys as operators emphasised $\sim 24°$C carriages/stations, relative to ambient temperatures of $\sim 50°$C, and public advisories urged avoiding mid-day exertion [341, 481, 485]. With others reporting a switch from walking to rickshaw to limit exposure [482], this is consistent with the theory that heat is a tax on mobility, in this case a monetary one, as it forces those who would walk onto transit. Informal services thin out on hot days: with multiple sources reporting a hawker survey indicates sharp declines in customers and $\sim 50\%$ income hits for Delhi street vendors; drivers and delivery workers report longer breaks, shading and cooling improvisations, and reduced mid-day availability [488, 227, 486]. Spending behaviour changes as people attempt to adapt—purchasing also air conditioners, but also cold drinks—which rearranges errands. Water scarcity in poorer colonies pushes reliance on tanker markets, raising the shadow cost of mid-day trips; the net effect is mid-day contraction with partial evening substitution and indoor/AC reallocation [402, 372, 43], consistent with what we find in our data.

Indonesia. Even when episodes do not meet the government's criteria for a "heatwave", official warnings about mid-day exposure, and elevated risk for outdoor workers; guidance repeatedly encourages shifting errands to mornings/evenings [238, 30, 365, 366, 160, 179]. Reporting emphasises adjustments in *timing* and *mode* rather than widespread closures (e.g., gig workers and informal services concentrating activity outside peak heat), which is consistent with smaller measured heat elasticities than in India or Mexico [366, 179].

Mexico. Record temperatures and grid alerts/blackouts accompany behavioral shifts: mid-day outdoor commerce and errands contract, with consumers substituting toward air-conditioned venues and online channels [403, 404, 184]. Papers of record and regional outlets document *ambulantes* reducing hours and reporting 50-60% sales drops as footfall evaporates in early afternoon; authorities shorten or suspend school days during peak heat [354, 155, 378, 364]. Scarcities of cooling goods—including ice rationing at one of Mexico's largest grocers—further raises the marginal cost of going out and weakens the business case for mid-day vending, reinforcing evening reallocation [368].

Box 6

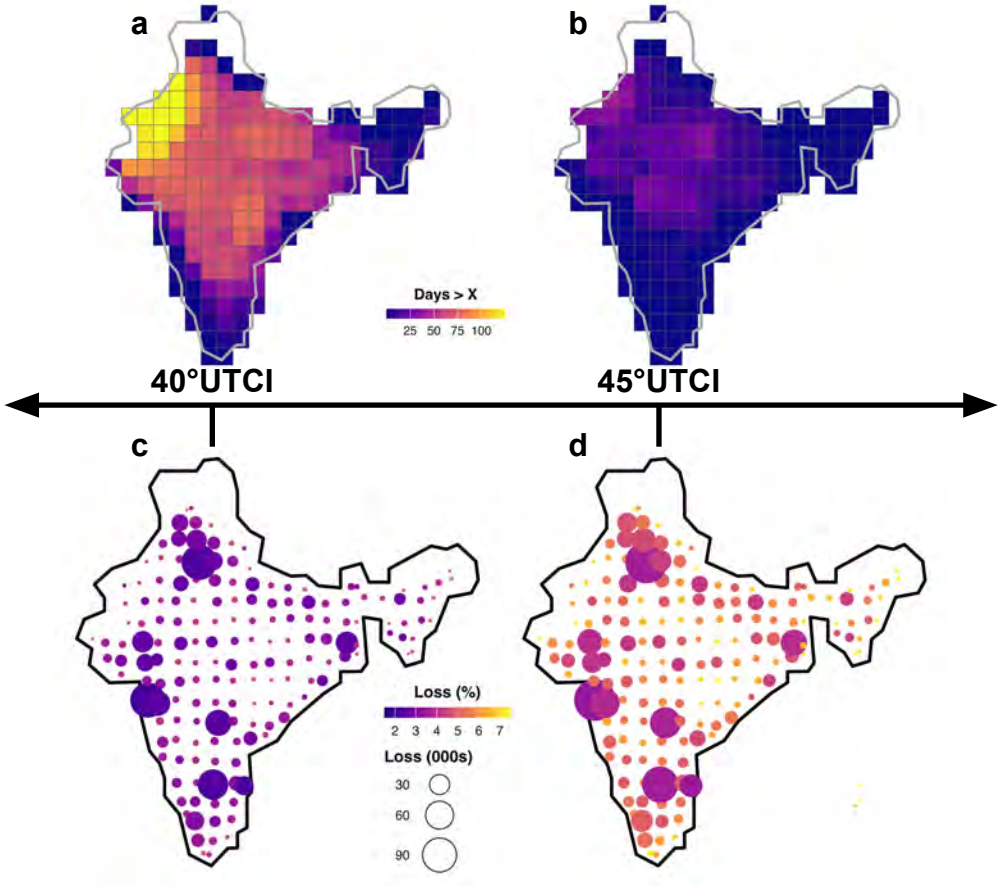

Figure 8.7. **Future implications for India. A** CMIP6 projections for days above 40°UTCI in India by 2050; these indicate that Northern India could experience as many as 150 days per year at or above this temperature. **B** CMIP6 projects fewer days above 45°UTCI but much of the North will still experience 25 days at this extreme. **C** and **D** show the consequences of this heat according to our modeling: during those days between $40 - 45^{\circ}$UTCI, we can expect activity to decrease by 5% in small cities and 2% in large agglomerations; when the temperature is higher, the loss of activity could be as much as 7% and 5% in small and large urban areas respectively.

Indonesia does not have cold weather, rising temperatures there are associated with an decline of 2% across the year. In Mexico, activity will fall as temperatures rise, by 3% on aggregate. This points to changing dynamics: we can expect heat waves to increase in both frequency and magnitude, and thus for associated disruptions to increase, but rising temperatures will have mixed effects. These effects will be spatially and temporally heterogeneous, affecting cities and seasons in different ways.

Discussion

Heat waves are expected to become a regular occurrence in urban areas over the coming decades. While efforts to mitigate climate change are important, we must also develop adaptive strategies. Understanding the relationship between human behavior and extreme weather, such as heat waves, is essential for this. Our study considerably expands to a growing body of research by examining how extreme heat affects human mobility in three large developing countries. Our findings show that heat waves significantly change urban mobility patterns. By examining three diverse counties across different climate zones, our research allows for comparisons and helps identify broader behavioral trends. The

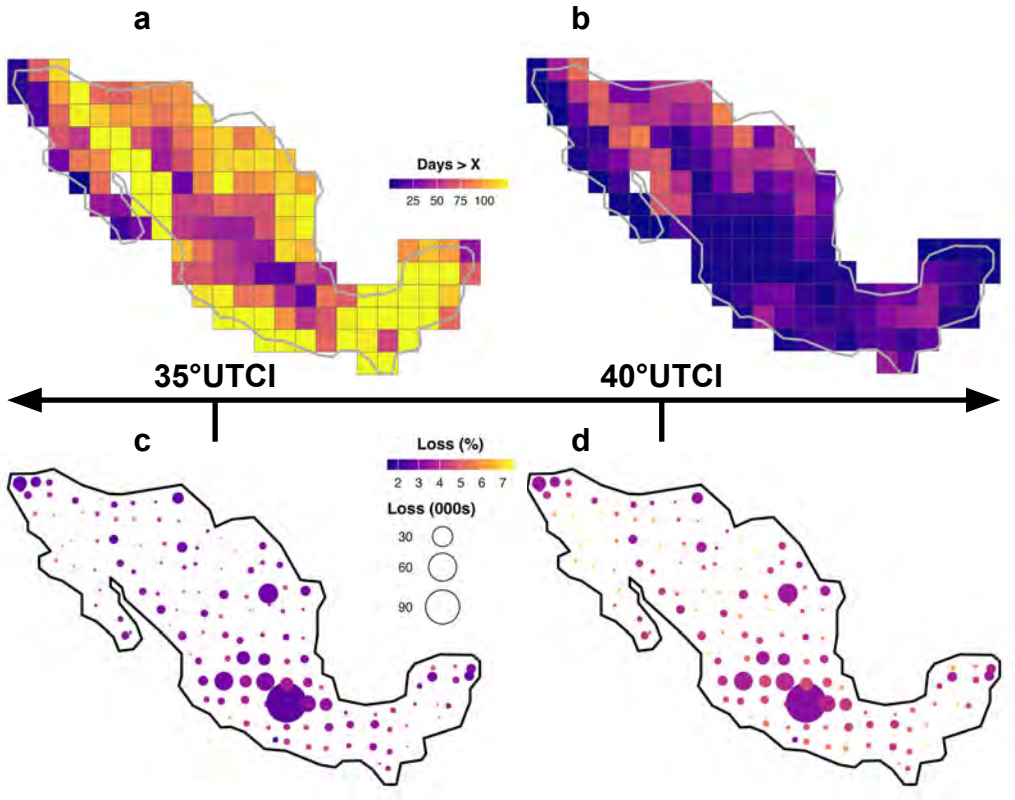

Figure 8.8. **Future implications for Mexico.** **A** CMIP6 projections for days above a given threshold in Mexico by 2050, **A** and **B** show frequency of days above 35°UTCI and 40°UTCI, respectively; we see that the tropical South will have more frequent days above 35°UTCI but the arid North will have more days above 40°UTCI, suggesting that heat ways will be a problem there. **C** and **D** show the consequences of these temperatures according to our model, with strong effects at higher temperatures but Mexico City, the capital and largest city, is spared of the worst effects because of its temperate climate.

highest activity levels and longest trips occur at the mean temperatures for each region. However, the specifics vary between these counties. For instance, in India, activity levels were more affected by extreme high temperatures than by lower ones, whereas in Mexico and Indonesia, the reduction in activities was almost symmetric between high and low temperature extremes.

Our results align with existing research on the economic impacts of high temperatures [460, 445, 532] which found that productivity losses are the result of lower crop yields or impaired worker performance; we expand on this by showing that extreme heat also influences travel decisions. For example, anecdotal evidence suggests that farm workers avoid outdoor labor when temperatures exceed 40°UTCI and even service workers need to change transportation modes to reduce heat stress from walking [134], and our data aligns with this while pointing to broader changes cities. In India, research on congestion indicates that people are willing and able to shift trips away from congested periods in the presence of taxation [262]; in showing that people substitute afternoon activity for morning activity, we find that heat functions as a tax on activity that people can avoid.

In India and Mexico, respectively, we observe a noticeable decline in activity levels when average temperatures exceed 33°UTCI and 35°UTCI. We use rolling averages to show that, in India, while longer heat waves have stronger effects, activity resumes when high temperatures persist for a week or more. While extreme heat never shows a positive

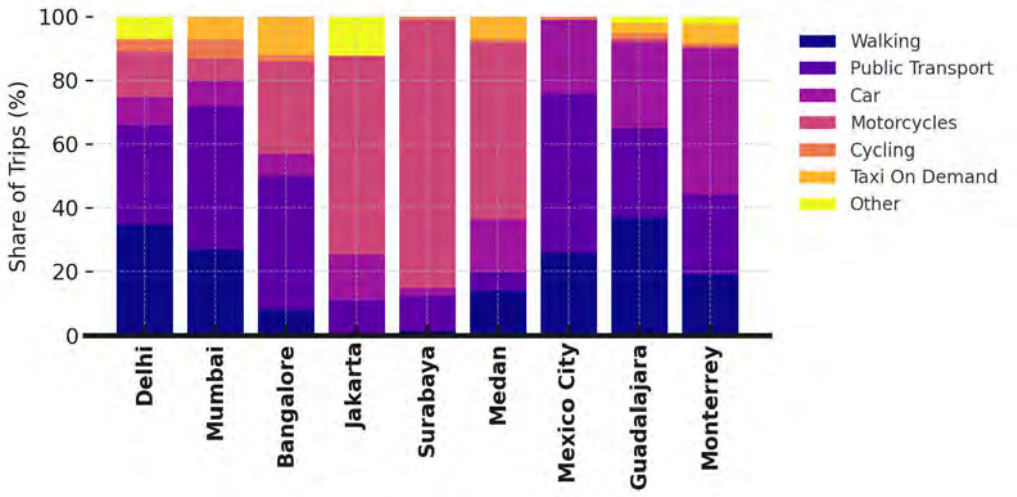

Figure 8.9. **Modes of transport in study cities.** Data from the Transformative Urban Mobility Initiative [491] and the Asia Transport Observatory [34] show that the dominant modes of transport varies in key big cities covered by our study. Motorcycles are common in Indonesian cities. Indian cities show more varied choices but walking is more prevalent. Mexican cities are wealthier on average and show higher car use.

effect on mobility in our data, the negative effect converges to 0 when the temperature averages more than $40°$UTCI for 6 or 7 days. We are unable to examine the mechanisms by which this occurs but if it is the result of individuals needing to return to work, it adds to the importance of cooling efforts.

However, the effects we observe are not equally distributed. They are more pronounced in deprived areas, indicating that these regions are more vulnerable to extreme weather disruptions. We find in India and Mexico that foot traffic in areas characterized by high deprivation fell faster at higher temperatures than foot traffic in those with low deprivation. Consistent with evidence on informal economies in India, mid-day losses concentrate in poorer areas, suggesting that heat disrupts neighborhood service loops that bind informal sellers to nearby buyers. When footfall falls, expected returns to opening a stall or traveling to wholesale drop below heat costs; buyers stay home, sellers follow—a thin-market feedback absent in Spain's formal, more teleworking core.

Changes in activity will result from combinations of forces: the ability of the wealthy to keep cool with air conditioning, for instance, might buoy activity in wealthy areas while pressures to earn might sustain activity in poor areas. Our findings suggest that for residents of deprived areas, the expected gains from working are less than the dangers from working at high temperatures. This is especially concerning because residents in these areas are more likely to hold jobs that require physical labor, making adaptation more difficult.

India shows evening substitution during heat waves—consistent with retiming activity rather than reducing it. In Mexico, a comparably less hot country on average, this adaptation appears to be lacking. While Mexico shows some recovery at night on hot days, the highest temperatures associated with fewer trips at all hours, pointing to stronger constraints from cooling access and infrastructure stress. We aggregate data on mode share in Fig. **8.9**. This is important when we consider who is most affected by high temperatures: for example, Indonesia's lower elasticities are interesting given its comparably higher use of motorcycles, which may enable more flexible travel.

The evidence we present here stands in contrast to what we saw earlier in Spain, where wealthier groups can avoid travel via telework and cooled residences, shifting the burden of adjustment away from subsistence work. Across settings, the same "tax on activity"

operates, but who pays it differs: in Spain, discretionary trips by the well-to-do or office workers; in India and Mexico, trips by the residents of deprived areas whose livelihoods may be dependent on hyperlocal economies, as informal work typically occurs close to home [5].

We should also note one key difference between Spain and the countries we study here: the size and structure of the informal economy. Informal work differs from contractual employment in that labor discipline is weakly mediated by formal contracts, allowing workers to make day-by-day cost–benefit decisions about exposure to heat. Two margins matter. Day labour (e.g., construction workers, loaders) is selected first-come, first-served; failing to show today does not typically jeopardise tomorrow's chances. Self-employment (street vending, home production) likewise decouples income from employer mandates, so opening hours and trips are adjusted to footfall and heat. Empirically, informality is pervasive in India ($\approx 90\%$ of workers), common in Mexico ($\approx 50\%$), and evident in Indonesia, with large home- and street-based segments that operate near home [241, 383, 108, 280, 33]. Spain, by contrast, has a smaller shadow/informal sector by both employment and value-added shares, reflecting stronger formal institutions [308]. Finally, cooling access is skewed: air conditioning and cooled travel are less prevalent among the poor in India, Indonesia, and Mexico than in Spain, amplifying the heat cost of mid-day trips: in Spain, transit usage is common amongst lower socio-economic strata, but in poorer countries it still represents a cost above walking [240, 151, 353].

In hot spells, informal work's margins of adjustment—opening, timing, mode choice—interact with heat's productivity and health effects. Heat reduces productivity and raises health risk in open-air, labor-intensive settings, decreasing the expected return to work and increasing the cost of travel; where revenues depend on street footfall, vendors may rationally shorten hours or delay trips [461, 108]. Because informal services cluster around home and rely on mutual trade [5], the buyer-stays-home/seller-stays-home loop becomes self-reinforcing at mid-day temperatures; evidence from Mexico City also shows that informality reduces the need for long commutes, concentrating activity in local loops that are more vulnerable to ambient heat [471, 108]. Spain's greater formality and higher feasibility of telework shift adjustment toward schedule changes by the wealthy than supply contractions by the poor, explaining the reversed socio-economic gradient in heat sensitivity we document across chapters.

A limitation in our study is the coverage of the mobile phone data that we use. The data are sparser in India than they are in Mexico, with Indonesia in between, and compensate for potential biases by filtering our data to cities, where coverage is best. In Mexico, where we are most confident in the data, we see the strongest effects and the cleanest differentiation across the temperature gradient in both the GAM and TWFE models. That the data in each country suggest an optimal temperature is also encouraging, but the future work will need to triangulate these findings. Yet the evidence for displaced activity in India warrants greater consideration as, if verified and validated with other datasets, it shows an important channel by which cities can adapt to climate change.

Finally, by modeling future climate scenarios, we estimate how human activity will be affected in the years to come. These projections allow us to better understand the evolving situation, not just due to changing weather patterns but also due to broader shifts, such as the planned relocation of Indonesia's capital from Java to the island of Borneo. Heat waves in India's Rajasthan and Indonesia's Sumatra will have the strongest effects on mobility, with temperatures exceeding 40°UTCI more than 100 days per year in both, changes that are associated with 10% declines in activity according to our models.

A remaining question is whether or not populations will gradually adapt with rising temperatures. Some places will have stronger heat waves, like Mexico's arid North, and

others will have higher but stabler temperatures, like its tropical South. Recent work on economic growth and public health looking across 5 decades suggests the answer is no: for most variables, responses to extreme heat have not changed [94]. An important finding in our work is that more deprived areas see stronger effects than less deprived ones, which has implications for adaptation: these areas will be amongst the last to adopt air conditioning. In the coming decades, our results may hold for vulnerable populations even if others are able to adapt.

Methods

We use origin-destination matrices for three nations—India, Indonesia and Mexico—spanning 2019. These data contain flows between tessellated units, called geohashes. The geohashing system creates a nested hierarchy wherein longer codes indicate sharper resolution. To preserve the anonymity of users, all data have been aggregated by the data provider to the geohash5 level—with cells that are on average $\sim 22\text{km}^2$—and do not include any individual records. Flows are summed from trips between any geohash6 cell where a device has recurring GPS signals [533], and can thus occur in the data within the same origin and destination geohash5 as a loop. Aggregated data were provided by Spectus Social Impact as part of the Netmob 2024 conference. Data are collected with the informed consent of anonymous users who have opted in to anonymized data collection for research purposes.

For each geohash5 in the data, using Google Earth Engine [209] we also gather 365 days of weather data from the ERA5 daily climate aggregates [133]. These data gives us temperature, in the form of Universal Thermal Climate Index (UTCI), per day per geohash5, and it allows us to observe differences within and between regions without relying on weather station data. We acquire modeled estimates of these variables for each day of 2050 using spatially disaggregated data from the Coupled Model Intercomparison Project [480], an international collaboration to produce estimates of Earth's climate through 2100. These datasets both resolve to $\sim 775\text{km}^2$, and thus also capture spatial heterogeneity in weather conditions. Finally, we attach data—via zonal statistics—on population [440] and deprivation [191] to each geohash5, using gridded measures estimated from remote sensing.

Using these data, we perform further checks on these data to understand is coverage and reliability in Appendix Fig. **C.8**; while the number of devices in a given cell shows a strong correlation with the number of residents in it, we also see that more deprived areas are less represented in the data. Most of this bias comes from cells with very low population, which are also typically poorer. Indeed, the correlation between deprivation, so we clip our data to "functional urban areas" [441]. (We use multiple filters, show in Table **C.5**, to stress test our models, but in the main analysis we use the "low" constraint filter.)

It is important to note that there is limited data in the context we are exploring in this paper, so we maintain that our work is still valuable even in the presence of bias.

Modeling the causal effect

We employ a two-way fixed effects (TWFE) approach to model the relationship between temperature anomalies and mobility in a given geohash5 cell, specified as

$$\ln(y_{it}) = \beta_0 + \beta_1 \Theta_{it} + \varphi_i + \nu_t + \varepsilon_{it} \tag{8.1}$$

where y_{it} represents mobility (measured as the number of trips ending in cell i on day t). The independent variable Θ_{it} captures temperature stress in cell i on day t, which we define as the deviation from the cell's mean temperature over the study period. This specification allows us to examine how departures from typical local conditions affect mobility patterns, rather than absolute temperature levels which could conflate adaptation effects.

The model includes two types of fixed effects: φ_i represents cell-specific fixed effects that control for time-invariant characteristics of each geographic unit, such as infrastructure, elevation, or proximity to water bodies. ν_t captures day fixed effects that account for temporal factors affecting all cells simultaneously, including holidays, weekends, and seasonal patterns. These day fixed effects also help control for potential sampling variations in the GPS data collection, as the number of active devices may fluctuate based on the set of applications providing data on any given day.

To account for potential heterogeneity in temperature effects across different times of day, we estimate separate models for morning (6:00-12:00), afternoon (12:00-18:00), and evening (18:00-24:00) periods. This temporal disaggregation reveals how mobility responses to heat stress vary throughout the day, capturing potential behavioral adaptations such as shifting activities to cooler hours.

We cluster standard errors at the cell level to account for potential serial correlation in mobility patterns within geographic units. The coefficient of interest, β_1, represents the semi-elasticity of mobility with respect to temperature stress, interpretable as the percentage change in mobility associated with a one-unit increase in temperature deviation from the local mean.

Modeling the temperature gradient

We begin using a generalized additive model (GAM) to measure the partial effect of temperature at different levels [514]. We build a directed acyclic graph (DAG) to determine what controls are necessary in our model, shown in Fig. **C.9A**. In addition to temperature, solar radiation and precipitation could also modulate activity, with rain dampening activity and sun heightening, so we control for these along with day-of-year. Day-of-year is modeled with a cyclic cubic spline, which captures seasonality by allowing for variation by day while forcing the value of the spline at the start to equal its value at the end. Because geography and area characteristics will also influence activity, we use various geohash fixed effects, allowing the intercept to vary according to the unique activity profile of each geographic unit. A key assumption with this approach is that the response curve is the same across space—only the intercept will vary. Together, these spatial and temporal controls constitute the minimally sufficient adjustment set. While they are not confounds, adjusting for day-of-week and holiday effects improves the precision of our estimates. The resulting equation is

$$T = \beta_0 + f_1(x_1) + f_2(x_2) + \mathbf{X}\boldsymbol{\beta} + \varepsilon \tag{8.2}$$

where T is the number of trips ending in the geohash5, $f_1(x_1)$ represents the smooth function of temperature, $f_2(x_2)$ represents the smooth function of day (with a cyclic cubic spline), $\mathbf{X}$ is a matrix of control variables, and $\boldsymbol{\beta}$ is a vector of coefficients for the control variables. ε is the error term.

Modeling the time series

Because GAMs do not have an explicit time component, we then model the time series as a robustness check. To analyze the relationship between activity and temperature while

accounting for temporal dependencies, we employ an Autoregressive Integrated Moving Average model with Exogenous variables (ARIMAX) [237]. This model combines autoregressive terms, moving average components, and exogenous factors to capture the complex dynamics of the time series. Specifically, our model incorporates AR(1) and AR(7) terms to account for immediate changes and day-of-the-week patterns in activity. The moving average component allows the model to consider the impact of past shocks and manages seasonality. Temperature is included as an exogenous factor with both linear and quadratic terms to capture potential nonlinear effects on activity. Our model is specified as

$$T_t = c + \phi_1 T_{t-1} + \phi_7 T_{t-7} + \beta_1 \text{Temp}_t + \beta_2 \text{Temp}_t^2 + \theta_1 \varepsilon_{t-1} + \varepsilon_t \tag{8.3}$$

where T_t represents economic activity measured as Trips at time t, $UTCI$ is a constant term, ϕ_1 and ϕ_7 are autoregressive coefficients, β_1 and β_2 are coefficients for linear and quadratic temperature effects, Temp_t is the temperature at time t, θ_1 is the moving average coefficient, and ε_t is the error term. This model structure allows us to simultaneously account for time-dependent patterns in economic activity and the influence of temperature, providing a comprehensive framework for analyzing the complex relationships within our data.

We compute the average coefficient for Temperature^2, which indicates the degree to which the effect of temperature curves at extremes, and we identify the ideal temperature by making predictions using the model for each time series at all temperatures and observing the point at which predictions begin to fall. We run a sensitivity analysis because not all time series span the full length of the year, showing the average coefficient at various minimum thresholds.

Projecting into the future

In order to understand how rising temperatures will impact behaviors, we take two approaches. First, we take the covariates from Eq. 8.2 but switch the temperatures on any given day for temperatures from the same day-of-year in 2050 using CMIP6 SSP 8.5 [480]. This gives us an aggregate estimate, incorporating both warmer winters (which might be good) and hotter summers (which might be bad). Second, we explore specific temperature thresholds that our modeling has shown to exert strong effects on mobility and explore how often these thresholds will be crossed according to CMIP6 SSP 8.5 estimates in 2050. To see estimates of temperatures in our 3 study countries through 2100, see Fig. **C.15**.

9 Conclusion

This thesis set out to measure how cities generate interactions between citizens, how they reconfigure when work detaches from place, and how they retime and redistribute activity as heat rises. We find that urban form drives experienced segregation, remote work reconfigures spending and mobility, and extreme heat reduces and retimes activity. The most important lesson from this work can be summarised as such: centres mix, peripheries sort, but remote work and heat threaten to change this dynamic.

Experienced segregation. Our first set of findings involves the structure of cities—the combined spatial distributions of amenities, activities, and populations—and how it produces experienced segregation. We first document strong regularities: American cities are typically surrounded by a ring isolating suburban neighbourhoods with residents who do not see people from the full range of socio-economic strata in daily life; these cities—especially denser, older cities—often have pockets of segregated urban neighbourhoods that do not host visitors from the broader city. These segregated pockets reveal that the effect of history on the present still lingers in mobility data, as they are often areas with complicated legacies of deprivation [512] and "hypersegregation" [301]—notably in South Los Angeles, South Chicago and South Bronx. Our results stand in contrast to studies that have emphasised the differences in clientele between restaurants or other amenities that are just metres apart [150, 321]. These studies miss something important: we find that whether or not an area is segregated, it is predictable using just a few key variables, like the median income of an area or its racial composition, twinned with its location relative to the centre. Yet earlier work that focused on variability over predictability holds relevant in centres and downtowns. Agreeing with this work, we also see that wealthy centres often have isolated populations that sort into homogeneous points of interest despite heterogeneous daytime populations, most notably in Manhattan but also in Chicago and elsewhere. Thus, predictability dominates at the mesoscale; variability appears within centres.

This is where predictability and variability interact in our work extending these results with a simple intervening opportunities model to understand how the location of amenities drives mixing. Although many residents of busy downtowns sort, likely because there is a surplus of options to facilitate sorting in these areas, centres are where amenities exceed expectations vis-a-vis social mixing. Central place theory, a key inspiration for the work in this thesis, explains why: it is the central places like downtowns that people are willing to travel farther to access. The intervening opportunities model assumes that a resident in a suburb should also dine out in that suburb, but if they are also anchored to an office downtown, the set of opportunities relocates.

Thus in this work we have also revised our understanding of bridging amenities. Prior work showed that bridging amenities are important to the broader mixing patterns in a city, specifically finding that larger cities have higher experienced segregation, but that the presence of amenities located between distinct socio-economic communities amelio-

rates this regularity. We extract central business districts and show that this scaling law does not hold for them: the city may become more experientially segregated with size, but not its downtown. Our results have important implications as companies around the world adopt remote work [53], enabling fewer and longer "supercommutes" [323], and diminishing the importance of central offices—which brings us to our next set of findings.

Remote work. We then show that remote work is having durable effects on mobility. Rather than a temporary shock followed by reversion, many cities—notably those that were "superstars" before the pandemic [218]—appear to have settled into a hybrid steady state: office anchors persist, but for smaller shares of the population. We find a marked size gradient in which cities have recovered from the pandemic and which cities have not: across a range of ordering strategies—using distance or density, for example—the largest cities are those least likely to have returned to 2019 mobility, and the cities that stand out most are older rustbelt cities like Washington, Philadelphia, and even Burlington. Newer sunbelt metros, which have grown in population and expanded in area over recent decades, have seen less change.

This reweighting matters for mixing. We see fewer trips between communities with different racial compositions—what we term convergent and divergent flows. Here again we see that the cities least likely to have recovered 2019 levels of these flows are larger cities and rustbelt cities. As we demonstrated in earlier work, when the office ceases to be an anchor, the intervening opportunities facing many residents relocate from the centre to the neighbourhood. Errands that were once chained to downtown workdays are now chained to home, shrinking the range of day-to-day activity and reducing the likelihood of intergroup exposure. Take together, our work suggests that both direct connections between different communities—convergent and divergent flows—*and* indirect connections moderated by points of interest in alternative locations like central business districts are threatened by remote work.

We show that the consequences of this reweighting away from cities and into suburbs have important implications for spending at local amenities. We find that the very places that we identified as rings of isolation in our first section are those that have seen spending recover since the pandemic. Again, though, the degree to which spending has shifted depends on the city—in particular where it is and how big it is. Expensive superstar cities like Los Angeles and New York have seen the Gini coefficient for spending—how spending is distributed throughout the city—rise since 2019; comparably affordable cities like Houston and Dallas have not. Our results suggest the most expensive cities in America were in a metastable state when the pandemic began. These cities were held together by intense agglomerative forces that remote work has attenuated.

Introducing a method that allows us to classify trips according to their lengths, we added an important caveat to this work on core-periphery dynamics and integration. We find that peripheral neighbourhoods are less efficient than central ones, committing to longer distances on aggregate—even after weighting on the number of trips. In light of research that suggests a fixed travel budget, this indicates that the automobile helps reclaim some of the time that would be lost to longer journeys, but it also shows that a world in which people must go to the central business district to work also forces residents of peripheral neighbourhoods to traverse larger portions of the urban extent on average than residents of central neighbourhoods.

Extreme heat. Using data on the present but with an eye toward the future, we show that heat influences mobility in both subtle and obvious ways. Generally, we see that on hot days around the world, people avoid going out during the hottest portion of the

day. Our results are intuitive, but we obtain point estimates that allow for more precise cost-benefit analysis. Using detailed data in Spain, we see that on balance, risks and incentives matter: older individuals reduce activity more on hot days than younger ones, commensurate with greater risk of illness at high temperatures; poorer individuals reduce activity less on hot days than wealthier ones, suggesting possible economic constraints. We also show that short trips, more likely to be taken by active means like walking or biking, are comparably more affected than longer trips more likely to be done by automobile—although we cannot identify travel mode in our data.

Expanding our analysis to India, Mexico and Indonesia, we find that the effects of heat are not uniform. In India, we see that activity falls during the day but rises at night on hot days, indicating what is called "intertemporal substitution", but we do not see evidence for this in Mexico. Our estimates are not statistically significant in Indonesia, which is an equatorial country with limited temperature variation, but they are still negative on hot days. An important difference that we see between developed and developing contexts, however, is that deprived areas show larger drops in activity on hot days in Mexico and India—where in Spain the poor appeared to be more constrained and did not change activity as much as the wealthy on hot days. Future work should explore this phenomenon in greater depth, but a plausible interpretation is institutional and sectoral: informal work is more prevalent in India and Mexico, and by design allows day-to-day adjustments in opening, timing, and travelling (self-employment, day-hire, piece-rate), while formal jobs in Spain exhibit greater rigidity. This raises the possibility that in large informal economies, hot days trigger a feedback loop—fewer people on the street, fewer reasons to set up shop, still fewer people on the street, and so forth. We do, however, see suggestive evidence for economic constraints in India and Mexico that fits with what we see in Spain: longer heatwaves show weaker effects over time, which hints at a possible need to restart work, errands and other chores as extreme heat drags on.

Finally, we see across countries, including Spain, that warm—but-not-hot days are good for activity, and we find suggestive evidence that a concept from public health—the "minimum mortality temperature"—inverts for urban dynamics: cities and regions appear to have a "maximum activity temperature" that hovers around the average temperature in that region. This idea also demands further research but it suggests that localities can adapt to rising temperatures—although it is not clear on what time scales these adaptations take place, since cultural and behavioural strategies to mitigate the effects of extreme heat could evolve over decades or centuries rather than years.

9.1 A bigger picture

We began this study by engaging with two contentions about the present and future of mobility. The first was that the past is foreign country, because they do things differently there. The second was that the future is already here, it is just not evenly distributed. These ideas frame the broader point that a reader should take away from this thesis. We motivated our work with the belief that cities are "social reactors" that generate wealth and innovation by pulling people together, allowing for both planned and spontaneous encounters, but this has been a study the forces that push people apart—from the emergent structures of cities, to new technologies and changing climatic conditions. New technologies and economic pressures mean that people move around cities with characteristically different trajectories than they did in 2019. Rising temperatures mean that the present in one part of the world might be the future in another. Combined, these forces will have important consequences for cities as social reactors.

The past is a different country. The urban present is not a reversion to 2019 from the pandemic. Prior to the COVID-19, superstar cities in particular were expensive; they were also crowded-yet-energised [458], facilitating face-to-face contact. Large technology companies that should have had the greatest capacity to virtualise were spending billions of dollars on vast headquarters busy downtowns: Google was expanding in New York City and London, Facebook and Apple in the Bay Area, and Amazon had announced two new offices in Washington and New York City. Today, remote work persists at large scales across countries, and the "doughnut effect" has dispersed activity away from expensive business districts—especially where the share of work-from-home days is high [12, 379]. Downtown recovery has been partial and uneven. In commercial real estate, the repricing of office buildings reflects a structural shift away from offices large enough to host all employees at a company [217]. Our work here builds evidence that sharpens our view of this new equilibrium: when day-to-day activity shifts to the suburbs, mixing falls in cities. The mechanisms are clear: remote work lowers the price of staying home, so the marginal lunch trip or coffee break with coworkers is less likely to happen [316].

Our findings should drive policy: a city centre that once mixed by default must now mix by design. This thesis shows how that reweighting trims intergroup exposure and weakens centre-to-centre links that normally deliver spillovers according to prevailing theories; the periphery sorts while the core—no longer guaranteed a large daytime influx of workers—must be animated to boost encounters at scale. This could take the form of programming and events, or a new mix of land uses. In short, while the claim that centres mix and peripheries sort holds across the work presented here, the very centrality of any given downtown is now in question. Central place theory holds that people will trade off the length and frequency of journeys to economise within fixed travel budgets; with fewer commutes, the geometry of travel budgets will change as people find themselves with more time. If central cities do not position themselves with these new travel budgets by offering attractions, activity will move to or remain in peripheral suburbs.

One reading of the evidence presented in this thesis is pessimistic: if distant trips and downtown anchors account for a disproportionate share of intergroup exposure, then the shift to remote and hybrid work cools the social reactor not only by removing the incidental encounters that cities once generated for free, but also by reducing planned face-to-face meetings. Yet what looks like a drawback can also be an opportunity to reorganise the city. Proposal for 15-minute cities or "urban villages" promise shorter distances and greater day-to-day autonomy, and our mixture results suggest that a proximate activity space is already remarkably stable across neighbourhoods even as higher-order travel reorganises. The risk is that proximity becomes a trap: dispersion can decrease distance while increasing isolation, yielding neighbourhoods that are walkable but socially separated. The alternative is a deliberate polycentricity: retain a strong, accessible nucleus for valuable matching and meeting and shared civic life, while cultivating diverse hubs and third places that draw residents across neighbourhood lines, supported by fast, affordable connections. In that framing, the decline of the monocentric commute does not have to mean the decline of urban encounter; it simply means that mixing must be produced—through design, programming, and connective infrastructure—rather than assumed.

The future is here. Some cities are already living in the future the rest will face more often. This thesis documents a common behavioural grammar of heat: activity compresses away from the hottest hours; routine and work trips are stickier than sporadic errands; and the largest absolute losses concentrate in big agglomerations even as smaller places suffer larger proportional contractions. The effects are mixed across developed

and developing countries, but economic constraints and incentives also play a role. In India, populations retime toward the evening during heat events; in Mexico, evening recovery is weaker, which implies adaptation is not costless—perhaps requiring solutions to coordination problems, since both buyers and sellers need to change timing together, and perhaps requiring expanded services, including transit.

Rank	Advice	n
1	Avoid drinking alcohol	34
2	Wear lightweight, loose fitting clothing	31
3	Drink regularly without waiting for thirst	28
4	Seek out an air-conditioned or cool environment	21
5	Stay indoors in an air-conditioned environment	20
6	Wear a hat	18
7	Avoid or reduce physical activities	18
8	Protect yourself from the sun	17
9	Know the symptoms of heat-associated illness, and know how to respond	15
10	Look in on susceptible people	15
11	Do not leave children in a closed, parked car	15
12	Avoid going out during the hottest part of the day	13
13	Take frequent baths or showers	11

Table 9.1. **Most commonly provided heat-protection advice to the general public**, with number of sources in parentheses [220].

This is what "the future is here" looks like in practice. Hotter places have already normalised delayed errands and staggered services; cooler places are beginning to show changes on the hottest weeks of the year, but fewer consistent adaptive strategies emerge. The portable lesson is not a one-for-one copying of habits but the enabling conditions: late (and safe) transit, shading and cooling measures, and extended hours will enable any country to become more like India or Indonesia. Where those supports exist, evening substitution preserves activity without commensurate health risk; where they do not, activity falls across the day and the feedback effects threaten informal workers. An uncomfortable possibility, however, is that the ways that we respond to extreme heat cut against what cities provide—social connection and "exuberant diversity" [242]. In Table 9.1 we show a collation of recommendations for coping with extreme heat [220], highlighting those actions that reduce urban vibrancy, like staying indoors, seeking air conditioning, or reducing activity outdoors during the hottest parts of the day. As days get hotter in many parts of the world, compressing the tolerable hours in a day, following these recommendations could mean that activity deferred becomes activity foregone.

We should note that learned adaption is not limited to heat: the shocks that we have examined in this thesis interact. Be they epidemiological, technological, or climatological, they demand renewed interest in combining social, natural and built infrastructure. Take parks as an example: the availability of green spaces is associated with reduced mortality from heat [328], improved mental health at the height of the COVID-19 pandemic [511], and, though recent evidence is still developing, helps mitigate some of the negative effects of screen time—notably by encouraging physical activity and improving mental health [110].

9.2 Limitations and future work

Our work here is limited along several key dimensions. We are often dealing with aggregate data that obscure types of activity and modes of transit. Exploring heterogeneities in both activity and mode will be important for building more precise estimates of the costs of extreme heat. It will also allow us to understand the threat of extreme heat to the social reactor: right now, we can only show that activity falls during the hottest parts of the day in Spain, and that this includes trips that end at work, but fewer lunch meetings and coffee breaks will be visible in more resolved data. Our work requires a logical leap—leaving and returning the office counts in our data as a trip to work, but so it would appear as fewer work trips in the afternoon—but we can build better models with better data. Our data are also limited by resolution in India, Mexico and Indonesia, where we aggregate to areal units that include large portions of any given city; understanding the internal dynamics on hot days will allow us to make better inferences. In particular, we have relied on triangulating our observations with other findings and regularities about the informal economy, but detailed data could shed light on specific changes to behaviour. Our data also require imputations: we infer socio-economic class in the US with administrative data and we classify areas in India, Mexico and Indonesia according to modelled estimates of deprivation.

Yet another limitation in our work comes from the problem that we have tried to solve with this thesis: mobility is not a stationary process; it is changing and many cities are still adapting to various transitions—including remote work and extreme heat. This means that our findings are subject to change. Future work should continue to compare and contrast the present with the past in order to better understand the future. Although the pandemic inspired many studies that construct time series to understand shocks to mobility, many of our classic models of human mobility—notably some "laws", or regularities that span countries or continents—are derived from data collected before the pandemic [17, 444].

Finally, this work was inspired by the ways that our urbanism creates prosperity from crowding and interpersonal contact. Having established the presence of new urban regimes, the next step is to revisit agglomeration and scaling. Theories of agglomeration rely on face-to-face contact, often using density as a proxy for human interaction. We have a unique opportunity to revisit our assumptions, as data on patents, wages, and other measures of agglomeration become available over time. Are cities today fit for purpose, and will they be once our urban structures adapt to the twin shocks of extreme heat and telecommuting, which change both the risks and rewards, respectively, to leaving the home, taking lunch with a friend or a client, and experiencing the city in all its messiness?

How mobility patterns will evolve in the coming years is not clear, but habits and preferences are malleable. In studies of rail strikes, looking at Deutsche Bahn and London Underground specifically, a significant fraction of riders who are forced into a new route will not return to the old commute; after the strikes abated in each city, ticket sales for buses between German cities remained 8% higher [64] while in London 5% of commuters continued to take the newer route [266]. In London, the subtle rethink that the strikes caused saved commuters 53 hours per year [266]; given the time saved, this implies that the cost or annoyance of experimenting with the commute for a single worker on a single day was $49. Prefiguring the pandemic, Hurricane Sandy disrupted transit in New York City, forcing some to adopt telework, and some of those people never went back to the status quo ante [251]. Surveys indicate that commuting is one of the least pleasant activities we perform daily [247, 256, 107], requiring in one estimation a 19% pay rise to compensate for the hit to wellbeing [470], which suggests that people

"satisfice"—accepting an available option as satisfactory—rather than optimise.

As a shock and an opportunity to rethink behaviours and habits, the pandemic was an example of a "moment of change" [510], and because unlike many of the disruptions above it was comprehensive in its reach—affecting government and citizens, employers and employees—it was also an opportunity to rethink policies and norms. The literature suggests that commuting is an activity that is both psychologically taxing—with a standard deviation increase in journey times corresponding to a decrease in happiness that represents 20% of the decrease from becoming unemployed [470]—and frequent enough to allow for experimentation, and yet many consider that experimentation to be such a hassle that they would need to be compensated $49 to do it [266]. Newer research shows that a job that allows remote or hybrid work can pay much less [143]—indicating that the pandemic caused many to assess the mental and temporal costs of commutes.

Our research indicates that while the moment of change that the pandemic represents may have passed, the push and pull between employees and employers, pricing in office and housing markets, and adaptations to extreme heat are still evolving. This gives researchers an opportunity to find points of leverage, find opportunities to improve cities—to not let a good crisis go to waste.

Appendix

A Experienced segregation

A.1 Data reliability

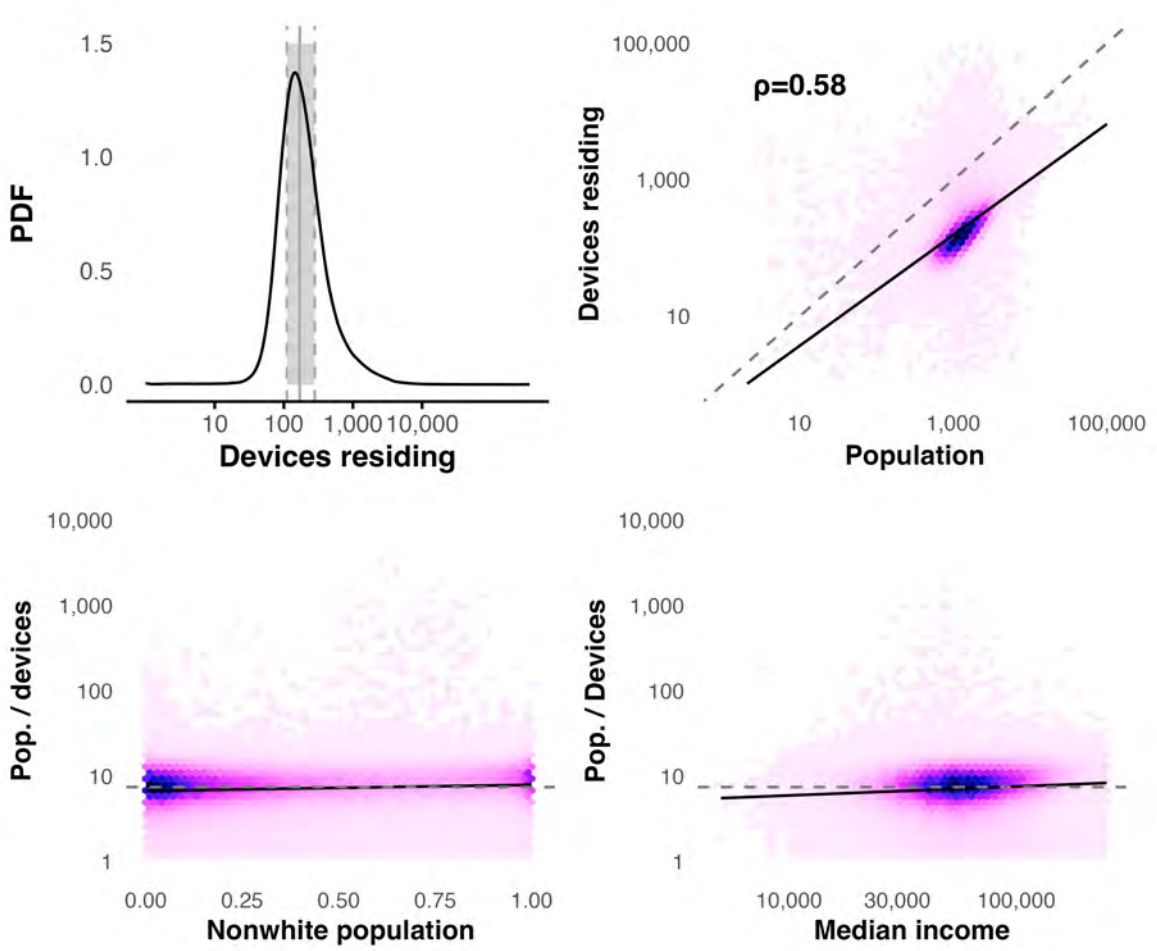

Figure A.1. **Data validation.** We test the reliability and representativeness of our data by first checking the distribution of devices: there is a tight distribution, with more than 150 devices in most block groups, and an interquartile range between 111 and 280. Further, we document a correlation between population and devices of $\rho = 0.58$, which is in line with previous studies, and more devices per block group that other studies [506]. Finally, there is no detectable bias with respect to income or nonwhite population.

Percentile	0%	1%	2%	3%	4%	5%	6%	7%	8%	9%
Devices	1	37	48	55	59	63	67	70	73	76
Percentile	10%	11%	12%	13%	14%	15%	16%	17%	18%	19%
Devices	78	81	83	86	88	90	92	94	96	98
Percentile	20%	21%	22%	23%	24%	25%	26%	27%	28%	29%
Devices	101	103	105	107	109	111	113	115	117	119
Percentile	30%	31%	32%	33%	34%	35%	36%	37%	38%	39%
Devices	121	123	126	128	130	132	134	137	139	141
Percentile	40%	41%	42%	43%	44%	45%	46%	47%	48%	49%
Devices	144	146	148	151	153	156	159	161	164	167
Percentile	50%	51%	52%	53%	54%	55%	56%	57%	58%	59%
Devices	170	173	176	179	182	185	188	192	195	199
Percentile	60%	61%	62%	63%	64%	65%	66%	67%	68%	69%
Devices	203	207	211	215	219	223	228	233	238	243
Percentile	70%	71%	72%	73%	74%	75%	76%	77%	78%	79%
Devices	249	254	260	267	273	280	288	296	305	314
Percentile	80%	81%	82%	83%	84%	85%	86%	87%	88%	89%
Devices	324	335	346	359	373	389	407	427	450	478
Percentile	90%	91%	92%	93%	94%	95%	96%	97%	98%	99%
Devices	510	547	594	654	730	833	979	1194	1541	2268

Table A.1. **Distribution of devices per block group.** The table shows the number of devices at each percentile of the distribution. The median block group has 170 devices. We only lose 2.5% of block groups when we limit to 50 devices.

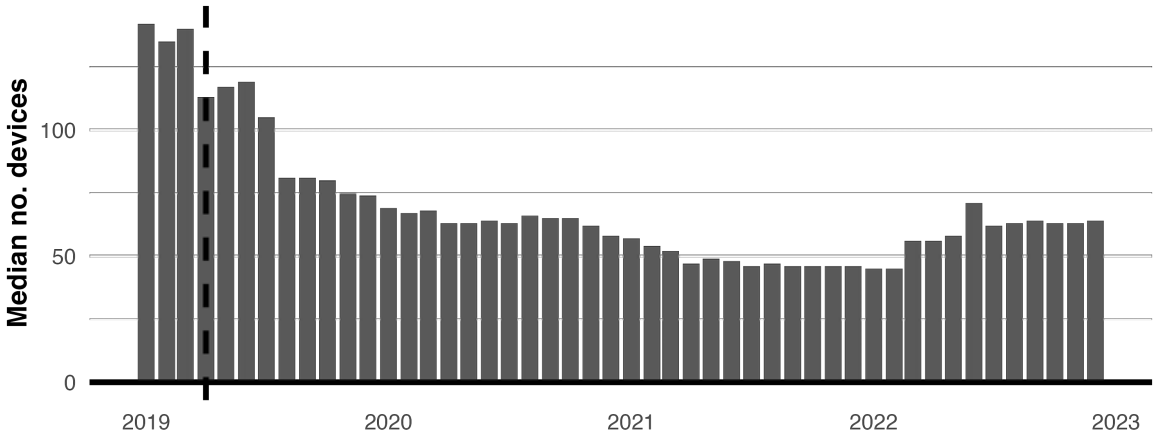

Figure A.2. **Device counters over time.** We show the median number of devices in each block group, documenting a dropoff after the first quarter of 2019.

Variable	**2019**	**2020**	**2021**	**2022**
Total Population	0.58 [0.58, 0.59]	0.67 [0.66, 0.67]	0.70 [0.70, 0.70]	0.66 [0.66, 0.66]
Nonwhite Population	0.02 [0.01, 0.02]	-0.07 [-0.08, -0.07]	-0.14 [-0.14, -0.14]	-0.10 [-0.10, -0.09]
Median Income	0.02 [0.01, 0.02]	0.11 [0.10, 0.11]	0.10 [0.10, 0.11]	0.06 [0.05, 0.06]

Table A.2. **Pearson correlations between device counts and demographic variables (2019-2022).** Each cell contains the correlation coefficient. Total population shows strong positive correlations with device counts in 2019 that grows stronger in later years, validating the device data as representative of population distribution. Nonwhite population and median income correlations are comparatively weak, suggesting minimal demographic bias in the device sample. Estimates are precise, with 95% confidence intervals (rounded) reported below.

A.2 The role of amenity mix

Distance quintile	Amenity count	Amenity entropy
1 (near to CBD)	0.038	0.042
2	-0.087	-0.101
3	-0.132	-0.139
4	-0.158	-0.164
5 (far from CBD)	-0.177	-0.144

Table A.3. **Spearman's correlation between median income and amenity mix.** The consistent negative correlation in higher distance quintiles shows how wealthy areas tend to have fewer amenities (amenity count) and less variation in amenities (amenity entropy). This pattern suggests an exclusionary mechanism where wealthier neighborhoods will see fewer visitors not through explicit barriers, but because of reduced place attractiveness—fewer points of interest and less diverse activity options that would otherwise draw visitors from different socio-economic backgrounds. (We use rank correlation because of skew in the data.)

Distance quintile	Segregation		Isolation	
	Amenity count (log)	Amenity entropy	Amenity count (log)	Amenity entropy
1 (near to CBD)	-0.426	-0.276	-0.009	-0.055
2	-0.357	-0.269	-0.076	-0.084
3	-0.333	-0.268	0.008	-0.027
4	-0.327	-0.277	0.047	0.007
5 (far from CBD)	-0.237	-0.080	0.392	0.013

Table A.4. **Pearson's correlation between amenity mix and segregation/isolation.** The amenity count shows strong negative correlations with segregation across all distance quintiles, providing robust evidence that areas with fewer amenities experience substantially higher segregation. This relationship holds across urban, suburban, and exurban areas, with correlations ranging from -0.33 to -0.43, until we get much farther from the core and the relationship falls off. For isolation, the correlations are negligible across most distance quintiles, contrasting with strong negative correlations observed for segregation. This difference suggests different mechanisms—and could be indicative of resident preferences in areas rich with amenities. (We use Pearson correlation because segregation and isolation are normally distributed and manage the skew in amenity counts with log-transformation.)

A.3 Bivariate classes and amenities

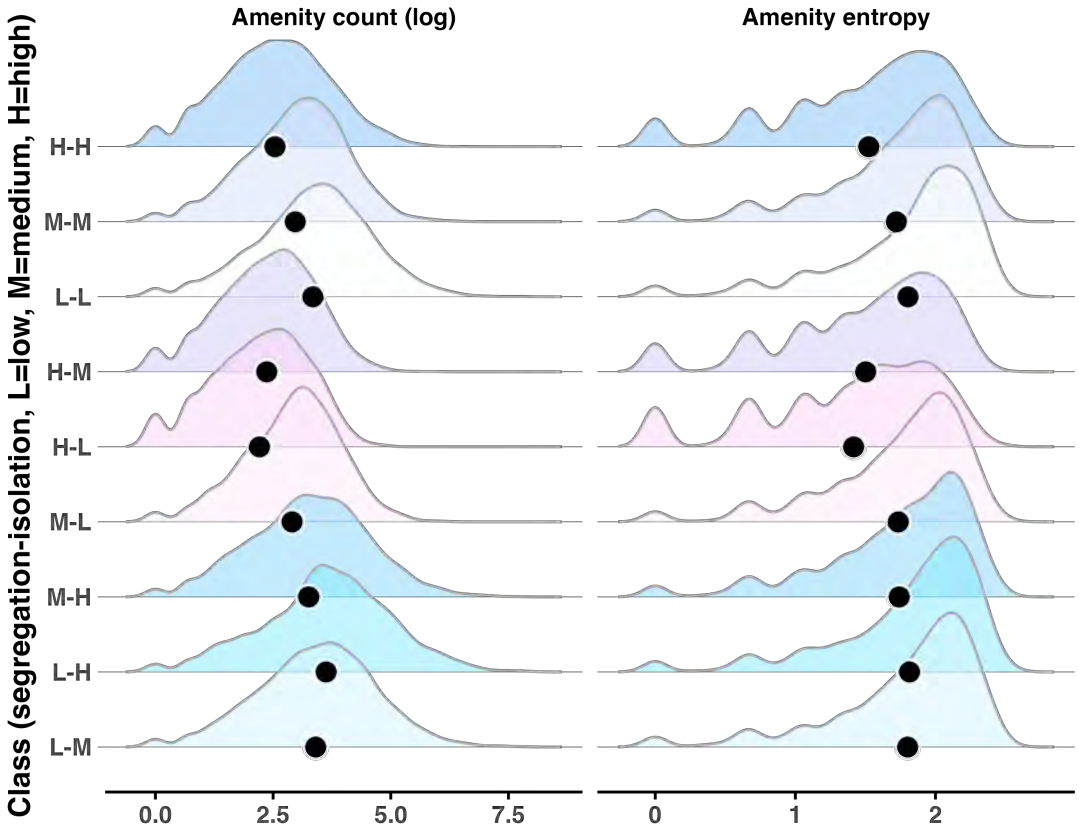

Figure A.3. **Amenity mix and each class in our schema.** While the relationship is mixed, we see consistent variation in how amenity count and the different classes of neighborhood relate: the bright pink "seekers" of diversity have the fewest amenities and the light blue "avoiders" of diversity have the most. Consider the diagonal, where isolation and segregation are aligned, fewer amenities begets higher S/I and vice versa. There is a similar pattern for the mix of amenities quantified with entropy.

A.4 Legacy of "Redlining" in our data

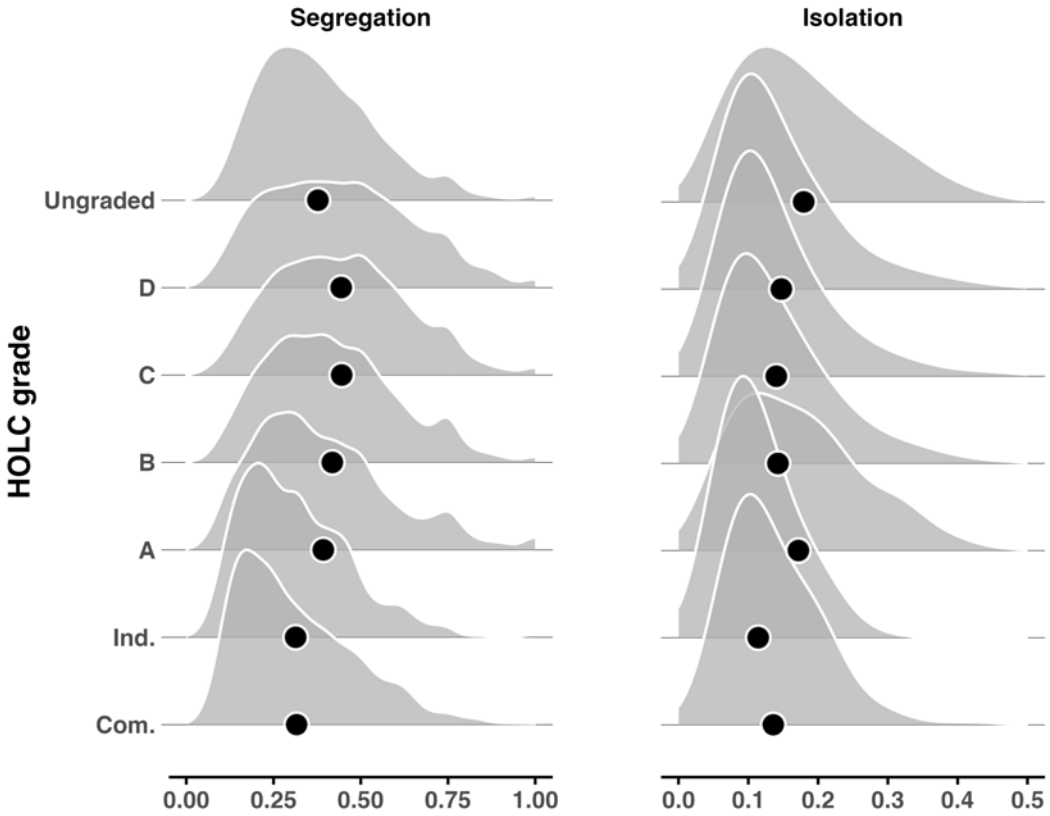

Figure A.4. **The legacy of Redlining in segregation/isolation.** Contemporary patterns of segregation/isolation vary with historical HOLC grades ("redlining" maps); distributions are shown with dots representing the mean value per grade. We see a clean gradient for *segregation* and a nosier one for *isolation*. Areas that were C/D graded, which often correspond to urban pockets, exhibit higher segregation, while ungraded areas—predominantly modern suburbs and exurbs that lay beyond the original HOLC boundaries—show higher isolation. This pattern reinforces our broader finding of isolated suburban rings surrounding cities with segregated urban pockets. The connectivity signatures in ungraded areas likely reflects the continuation of discriminatory practices through different mechanisms—while urban neighborhoods faced redlining, suburbs employed exclusionary zoning and other discriminatory practices [424] to maintain social and racial homogeneity. These historical practices appear to have created durable patterns in how communities interact, visible in contemporary mobility.

A.5 Histories of segregation in pockets

South Chicago. During the Great Migration, tens of thousands of African Americans moved from the Jim Crow South to Chicago, profoundly reshaping the South Side. New Black arrivals in the 1910s–1940s were met with a "different form of public denial" in northern cities: racially restrictive covenants, redlining by federal agencies, and discriminatory real estate practices that cordoned Black Chicagoans into the city's expanding "Black Belt" on the South Side. By the 1930s, more than 200,000 African Americans lived in this segregated South Side enclave, also known as "Bronzeville" or the "Black Metropolis," which became a center of Black culture and commerce despite overcrowded housing and limited opportunities. These structural barriers—enforced by law, violence, and federal housing policy—ensured that segregation "lasted through the 20th century and continues to define the city's geography" in Chicago, leaving the South Side predominantly Black and often underserved. [213, 477]

South Central. South Central Los Angeles emerged as a major Black community during the Second Great Migration, when African Americans flocked to California for defense industry jobs in the 1940s. Because racially restrictive covenants barred Black residents from most of Los Angeles, about 70% of the city's Black population by 1940 was packed along the Central Avenue corridor—the area that came to be called "South Central". World War II accelerated this influx: between 1942 and 1945 roughly 200,000 Black Americans migrated to Los Angeles, doubling the local Black population by decade's end. Yet Los Angeles authorities and homeowners reinforced segregation—housing covenants were enforced even more tightly as the Black population grew, so that the community "remained confined to pre-war boundaries" despite severe overcrowding. These practices created a segregated, majority-Black South Central enclave. In the postwar decades, as whites decamped to suburbs, South Central suffered from disinvestment, setting the stage for the area's persistent economic challenges and its later reputation for unrest. [464]

South Bronx. The South Bronx's postwar decline was largely the result of deliberate urban planning and institutionalized racism. The construction of the Cross Bronx Expressway (built 1948–1963 under Robert Moses) tore through stable Bronx neighborhoods, displacing 40,000–60,000 residents and leaving a gash of concrete that devastated the local housing market. Property values plummeted, and many white residents fled en masse to suburban areas (aided by government-subsidized mortgages), while Black and Puerto Rican New Yorkers often had no choice but to stay behind in the Bronx due to redlining and discriminatory housing policies. By the late 1960s the South Bronx was overwhelmingly populated by low-income African American and Latino families, and city services and private investment had largely dried up. This disinvestment culminated in the 1970s "Bronx is burning" era, when landlords torched buildings for insurance and entire blocks were reduced to ruins—roughly 80% of the area's housing stock was lost to arson and abandonment. These structural forces produced a South Bronx synonymous with urban blight and concentrated poverty, a legacy that continues to shape its social structure today. [474]

North Philadelphia. North Philadelphia became a segregated Black community during the 20th century through migration and systemic discrimination. In the 1920s and 1930s, Philadelphia's Black population swelled as Southern Black migrants arrived, roughly doubling the number of Black residents in North Philly and concentrating them in a few dense neighbor-

hoods. White Philadelphians increasingly refused Black neighbors—through racist violence, "neighborhood improvement" associations, and exclusionary real estate practices—leading to what researchers call hypersegregation, as Black families were confined to crowded blocks of North Philadelphia while whites moved elsewhere. By the mid-20th century, federal housing policies like redlining reinforced this pattern: virtually all of North Philly was deemed high-risk and denied mortgages, and nearly every public housing project built in Philadelphia between the 1940s and 1960s was located in these redlined Black neighborhoods. These forces entrenched poverty and blight in North Philadelphia's segregated communities—effects still visible today in the area's abandoned houses and chronic disinvestment. [187, 77]

Roxbury. Roxbury's modern form was shaped by deliberate segregation and aggressive "urban renewal." In the 1930s, federal Home Owners' Loan Corporation maps labeled all of Roxbury as "hazardous" for lenders—literally coloring the entire area red—because of its high Black and immigrant population. This redlining made home loans nearly unobtainable in Roxbury and reinforced a racial dividing line: after World War II, white ethnic residents left for the suburbs in a wave of white flight, and Roxbury became almost exclusively inhabited by African Americans (along with some Caribbean immigrants). Then, in the 1960s, Boston officials targeted Roxbury for highway construction and redevelopment schemes. Under the Boston Redevelopment Authority, nearly half of Roxbury's buildings were razed—ostensibly for a planned I-95 freeway and other projects—displacing over 5,000 families before community protests halted further demolition. These structural interventions devastated Roxbury's housing and economy, leaving behind a predominantly Black community cut off from Boston's prosperity by decades of disinvestment and physical isolation. [272, 473]

Anacostia. Anacostia underwent a rapid racial turnover in the mid-20th century due to desegregation, federal intervention, and white flight. In 1950, the neighborhood's population was 82% white and 18% Black (with Black residents largely relegated to a small enclave). After the Supreme Court's 1954 school desegregation rulings (including D.C.'s Bolling v. Sharpe), White residents initially protested but then fled en masse to Maryland suburbs rather than integrate local schools. By 1967 only 37% of Anacostia's residents were white, and following the 1968 riots that number dwindled to almost none. The departing White families were replaced by Black families—including many who were uprooted by a 1950s urban renewal project that cleared a large African American enclave in Southwest D.C., forcing those residents to relocate to places like Anacostia. These structural forces transformed Anacostia into a predominantly Black neighborhood marked by economic divestment after white flight, a legacy that continues to shape its social conditions. [156]

A.6 Sensitivity analysis

Census block groups, while the smallest geographical unit with available income data, contain heterogeneous populations. This internal variance raises concerns about whether our findings might be artifacts of using median income as our primary assignment method. To address this concern, we tested alternative income assignment approaches that capture different aspects of within-block-group distributions, allowing us to examine whether our observed patterns of socioeconomic mixing persist regardless of how we characterize neighborhood income composition.

To validate the robustness of our findings to the income assignment methodology, we conducted a sensitivity analysis using alternative approaches to assign socioeconomic status to visitors. In our main analysis, we assigned each visitor the median income of their home block group. Here, we test whether our measures of segregation (S) and isolation (I) are sensitive to this choice by recalculating them using three alternative income assignment methods:

- **Mean Income:** We assigned each visitor the average income of their home block group, which is more sensitive to outliers than the median.
- **Modal Income Bracket:** We assigned each visitor the most common income bracket (approximately $10,000 bins) reported in the Census for their home block group.
- **"High Variance" Scenario:** To simulate potential sampling bias in the mobile phone data, we deliberately assigned incomes drawn from non-modal brackets within the home block group. This represents a scenario where mobile phone users systematically differ from the typical resident of their block group.

Our results demonstrate strong robustness to these alternative specifications. The measures of segregation and isolation calculated using these alternative methods remain highly correlated with our original median-based measures, with correlation coefficients $\rho > 0.78$ for segregation/diversity and $\rho > 0.9$ for isolation/exposure (Fig. A.5). Furthermore, the distinctive spatial patterns identified in our main analysis persist across all specifications. The rings of isolation surrounding urban centers and the pockets of segregation remain evident in all alternative scenarios, with the patterns becoming even more pronounced in some regions under the high variance scenario (Fig. A.6).

To quantify this consistency, we examined how block groups were classified in our 3×3 isolation-segregation typology across different specifications. Approximately 80% of block groups remained in the same classification category regardless of the income assignment method used (Figure A.7). This demonstrates that our findings represent robust socioeconomic mixing patterns rather than artifacts of our income assignment methodology or potential biases in mobile phone representation.

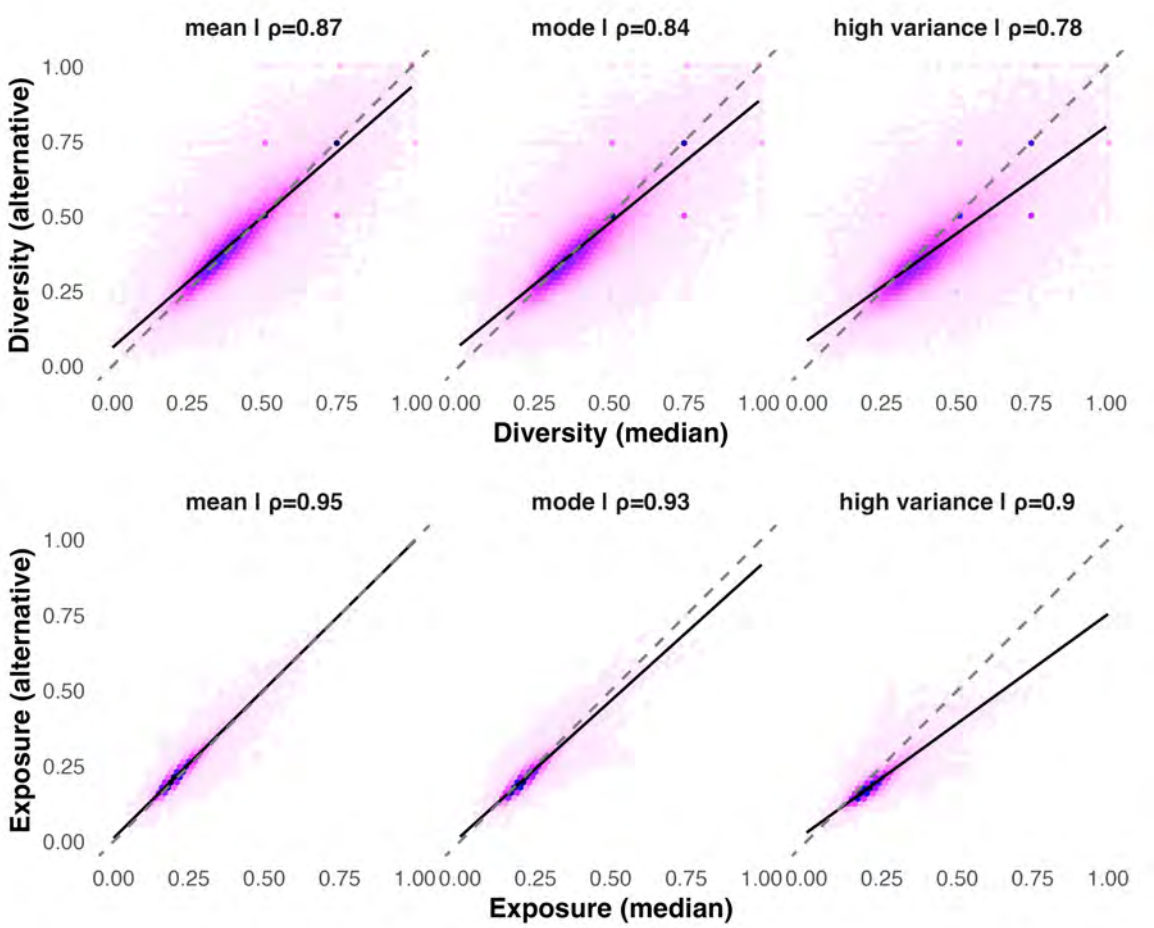

Figure A.5. **Sensitivity analysis.** We use alternative imputations that respond in different ways to underlying distributions and outlier behaviors. For example, we use mean income instead of median income—assigning each user the income of the average in block group. We also use the most common income according to the Census' income bins, which are $\sim \$10,000$, in the modal specification, as well as a scenario where we sample from the distribution holding out the modal bin. This represents a case in which mobile phone data are maximally biased, representing a population that is not the most common in block group. When we compute our measures of segregation and isolation using these alternative specifications, they show a high correlation with our preferred specification.

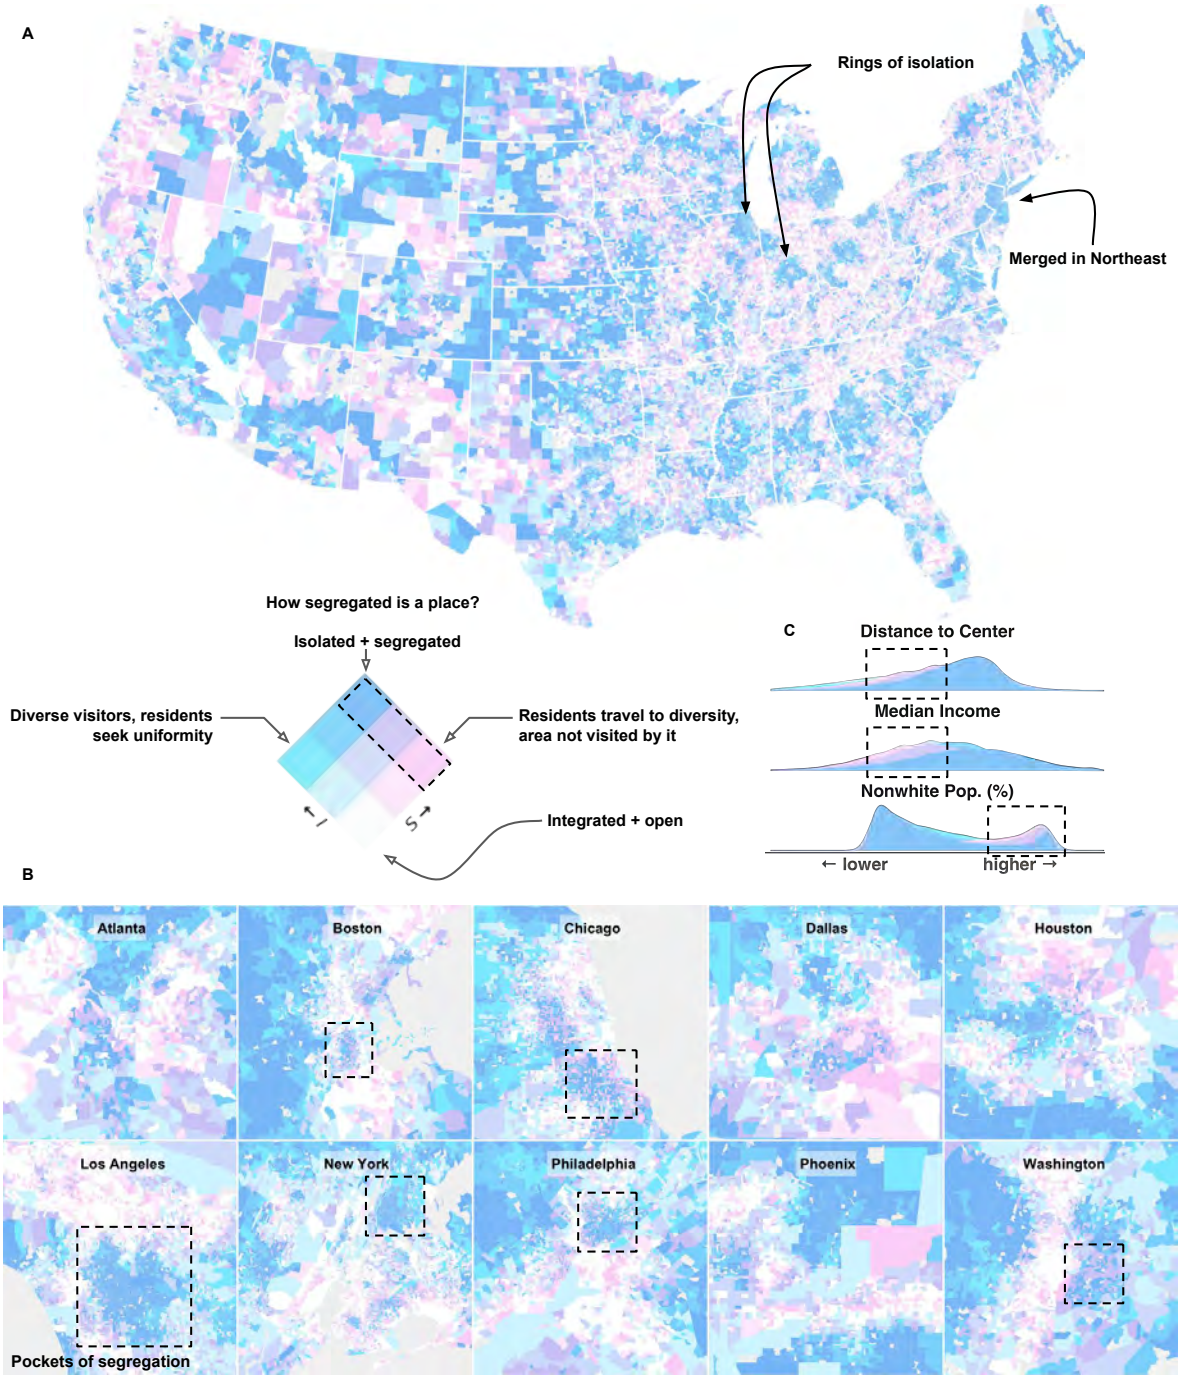

Figure A.6. **Ring/pocket robustness.** We test the "high variance" mixing scenario, showing that the results are indeed different but that the rings and pockets that still exist and in some cases are sharper and more defined.

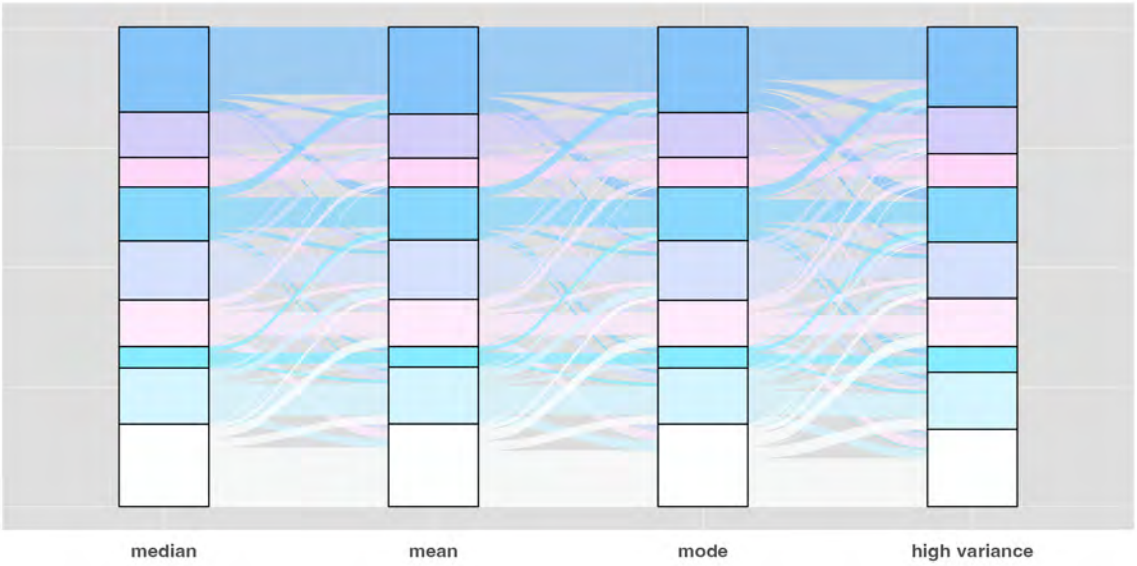

Figure A.7. **Transitions between classes with different income assignment strategies.** When we create our isolation-segregation classes using these alternative specification, most block groups are in the same class that they would have in our preferred specification with median income. Generally, 80% of all block groups are in the same class across all variations of our measures.

A.7 Are small CBGs driving our results?

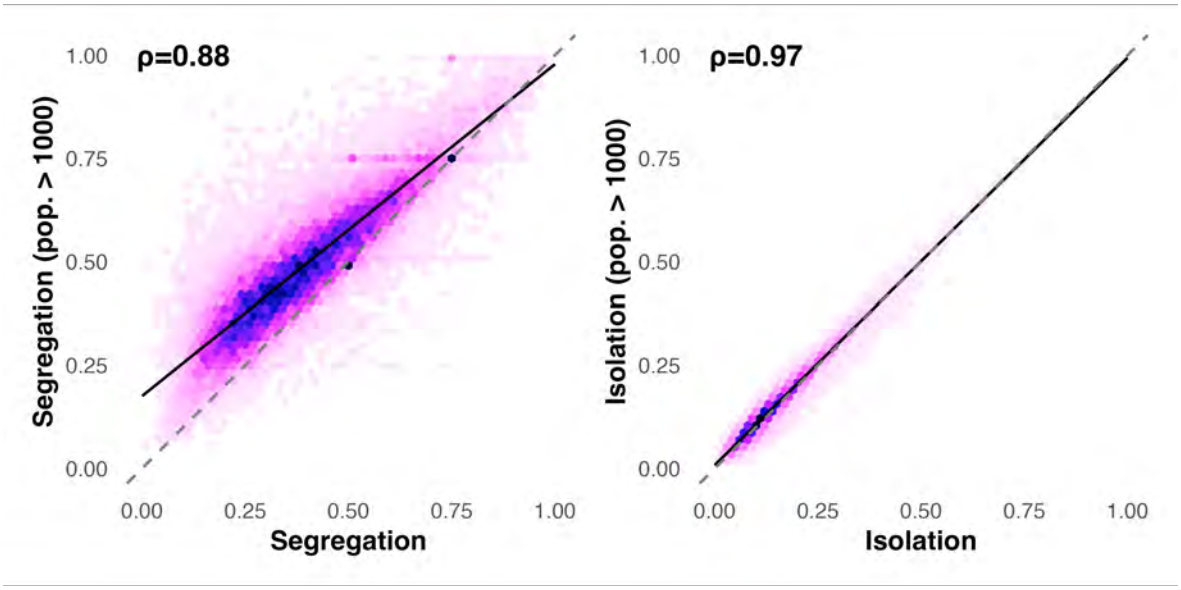

Figure A.8. **Sensitivity to population.**We ensure that our results are not being drive by rural areas with little population by filtering out block groups with less than 1000 residents and recomputing our metrics. This means that for Segregation (S) we only count visitors from these populous block groups and for Isolation (I) we use the new S value to compute exposure *and* only do so for block groups with the minimum population. Our results appear robust and the measures before and after filtering show high correlation.

A.8 Comparing income to nonwhite share

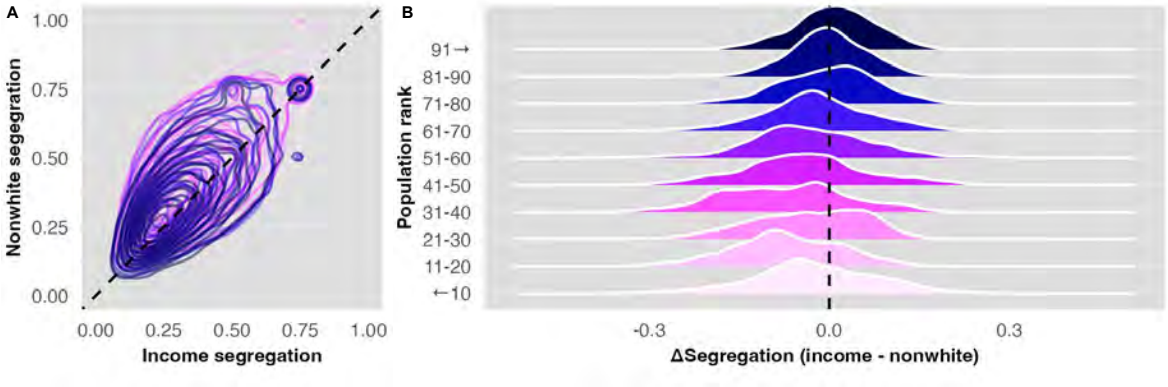

Figure A.9. **Comparison with race.** As a robustness check, we use estimates of the nonwhite population to check whether the relationships we see for segregation and isolation hold along dimensions other than income. In **A**, we show that place segregation using median income and place segregation using nonwhite population are correlated. In **B**, we subtract nonwhite segregation from income segregation for all observations and plot the distributions for different city sizes. For large cities, nonwhite segregation tends to be higher than income segregation, which strengthens our results. Comparison of segregation by estimated income and nonwhite population shows that although the measures are correlated, for larger cities by population, race segregation is systematically higher than class segregation.

A.9 Further robustness checks

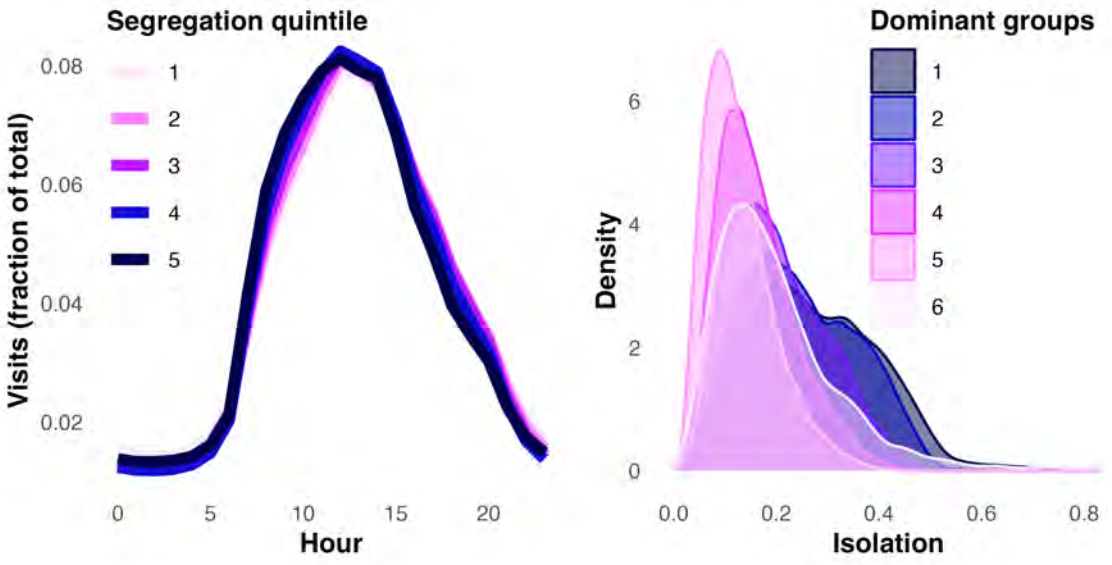

Figure A.10. **Time of day and composition effects.** Left: we partition amenities according to segregation quantile, then plot visits by time of day, seeing now differences across quantile. While we cannot explore whether or not segregation is different at different times of day, the consistency across all quantiles alleviates these concerns: if time of day corresponds with different visitors, we would see variation in these distributions. Right: we assign a "dominant group" to each amenity, according to its most common visitor income bracket, and then look at how many dominant groups different neighborhoods experience. The most isolate neighborhoods experience the fewest dominant groups, alleviating concerns that an isolated neighborhood could have a high value on this measure if its residents visit homogeneous, low-income amenities and homogeneous, high-income amenities (both would be high segregation). This does not seem to happen often in practice.

A.10 NULL model

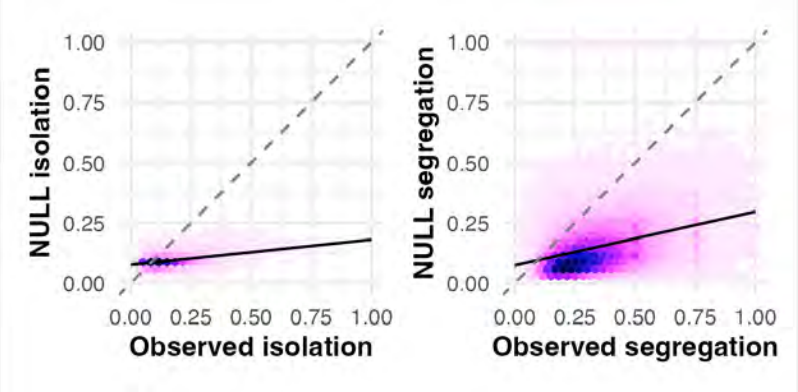

Figure A.11. **Counterfactual if people mixed at where they work.** Mixing compared to a null model. We construct our measures of segregation and isolation using commuting data acquired from the Census Bureau's LEHD Origin-Destination Employment Statistics [499]. We compute them for block groups and thus assume that all employees within one block group are exposed to each other. Segregation and isolation obtained from GPS data are systematically higher than when they are obtained from commuting data under these assumptions. This suggests that people sort into amenities around where they work, rather than mixing with the general population of workers. As the plot of isolation shows in particular, the assumption of general mixing is generous and there is less variation in these null results than in the GPS data, but it serves as a comparison to an unbiased counterfactual where individuals do not sort into amenities and simply mix at home and at work.

A.11 Features by importance

Rank	Variable	Importance
1	Median income	0.2589
2	Dist. to CBD	0.1543
3	Density	0.1445
4	Amenity (#)	0.1181
5	Nonwhite (%)	0.0991
6	College (%)	0.0799
7	Amenity (H)	0.0499
8	Household size	0.0319
9	Vacancy rate	0.0242
10	Rent burden	0.0195
11	Under 16 (%)	0.0126
12	Unemployed (%)	0.0071

Table A.5. **Table of important features in our tree.** Income is the single best predictor of segregation but urban structural factors like distance to the CBD, density and the number of amenities are also important. We also note here that, although it is collinear with income, unemployment is irrelevant in this model. We might expect, insofar as many jobs require travel to work, employment would represent a connection to the broader economy.

A.12 Complete decision tree

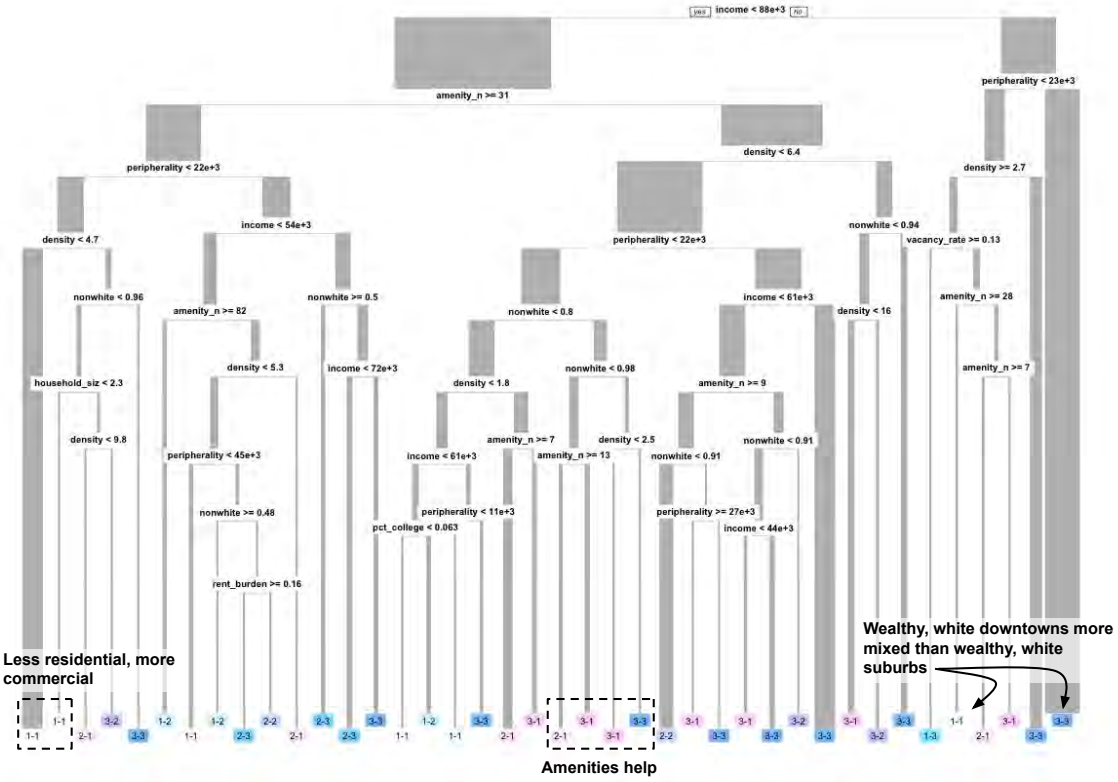

Figure A.12. **Full decision tree.** Decision tree with all predictors, showing that wealth, race, and location are dominant predictors, but also that amenities play an important role. There are segments of the population that avoid diversity, despite considerable variety in visitors nearby, but these groups are idiosyncratic, with a variety of features—that could be proxies for other facets of urban life—predicting them.

A.13 Relationship between spatial and social mobility

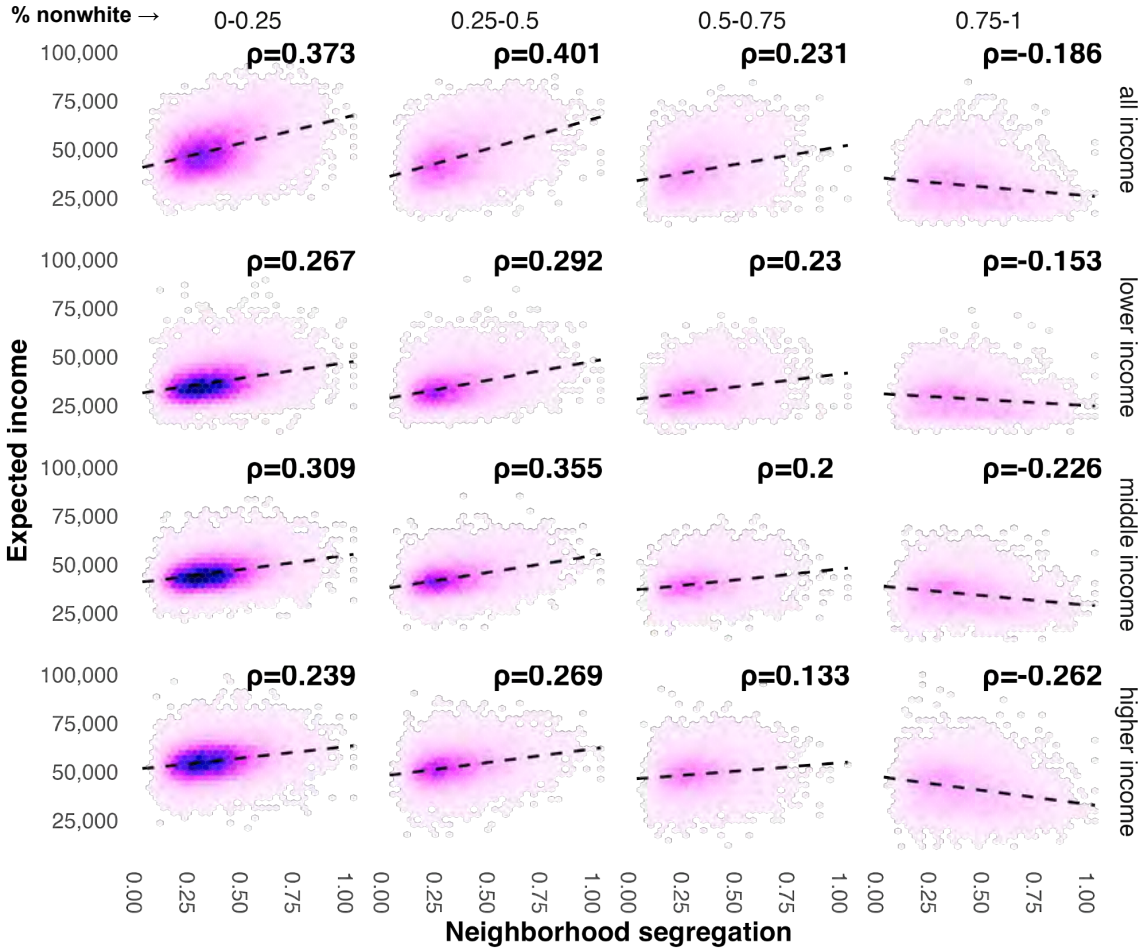

Figure A.13. **Social mobility and segregation.**Relationship between neighborhood segregation and expected adult income, across different neighborhood racial compositions and childhood income levels. In predominantly white neighborhoods (0-50% nonwhite), higher segregation paradoxically predicts higher incomes ($\rho > 0.3$), likely showing that it signals affluence. This relationship weakens as neighborhood nonwhite share increases and becomes negative in high nonwhite areas (75-100%), suggesting that visitor homogeneity may proxy for disadvantage in these communities, though the relationship cannot be interpreted causally.

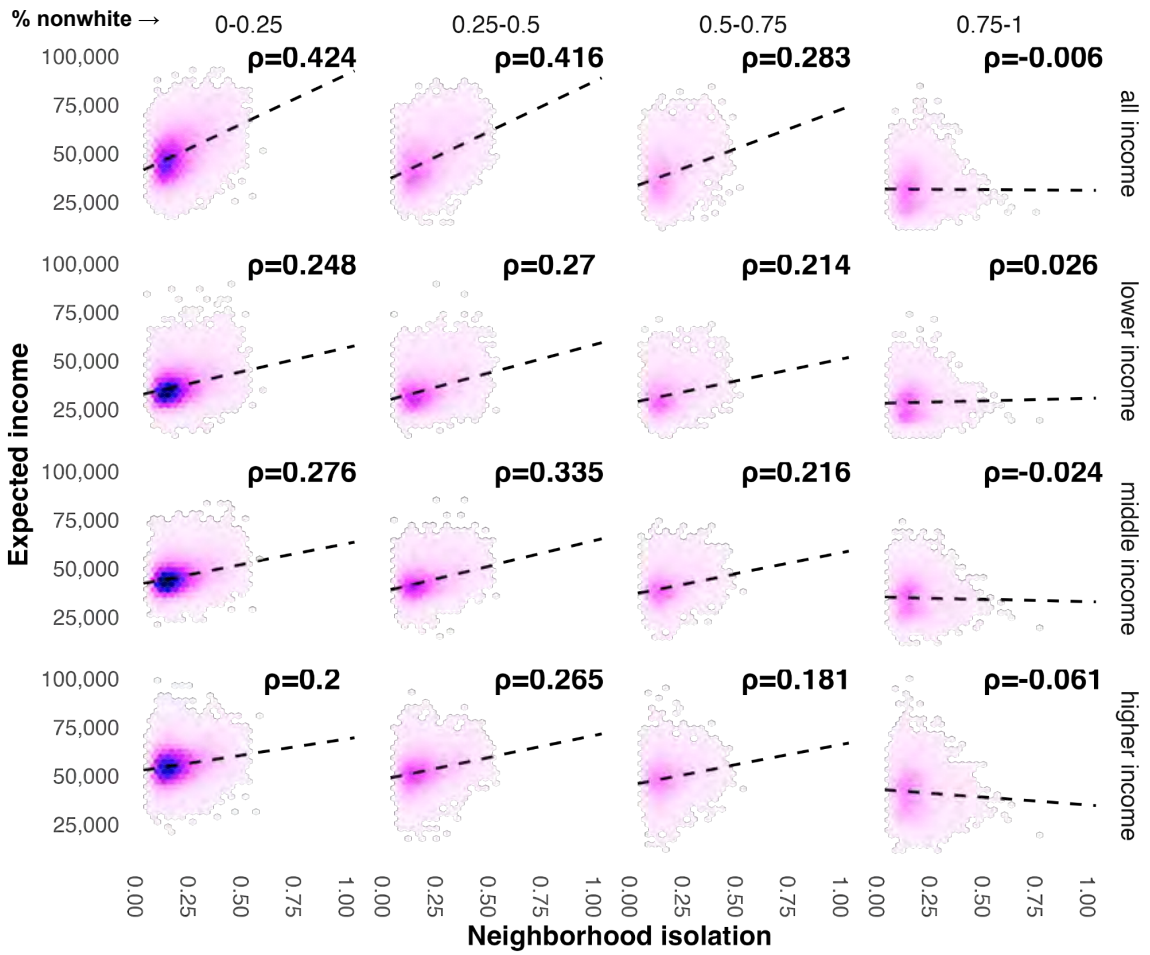

Figure A.14. **Social mobility and isolation.** Relationship between neighborhood isolation (homogeneity of residents' destinations) and expected adult income, across different neighborhood racial compositions and childhood income levels. Similar to segregation, isolation shows strong positive correlations with income in whiter neighborhoods ($\rho > 0.4$ for 0-25% nonwhite), again reflecting affluence rather than causal effects. However, unlike segregation, the relationship approaches zero rather than becoming negative in high non-white areas, suggesting that the diversity of visited places may matter less than the diversity of visitors for predominantly nonwhite neighborhoods. Because this seems implausible, it may again indicate that amenity segregation is a proxy for other factors that determine the economic opportunity in a neighborhood.

Variable	Adult income (SD)
Isolation (SD)	0.093***
	(0.017)
Segregation (SD)	-0.026*
	(0.011)
Median Income (000s)	0.017***
	(< 0.001)
Nonwhite pop. (%)	-1.412***
	(0.063)
Metro Area FE	Yes
Observations	70,687
Adjusted R^2	0.743
Within R^2	0.695
RMSE	0.504

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A.6. **Neighborhood connectivity and adult income outcomes.** Fixed effects regression results examining the relationship between neighborhood social connectivity and adult income. A one standard deviation increase in *isolation* corresponds to a 0.093 SD increase in adult income (approximately $1,274 annually), while controlling for neighborhood characteristics and metro area fixed effects. This suggests that isolation is a proxy for other aspects of a neighborhood, like the affluence we identify in rings. The *segregation* shows a small negative association (-0.026 SD), indicating that it does not explain much of the difference in outcomes in pockets. Traditional neighborhood characteristics remain stronger predictors: a one percentage point increase in nonwhite population share is associated with a 1.41 SD decrease in adult income. Standard errors (in parentheses) are clustered at the metro area level.

A.14 Analysis of “churn"

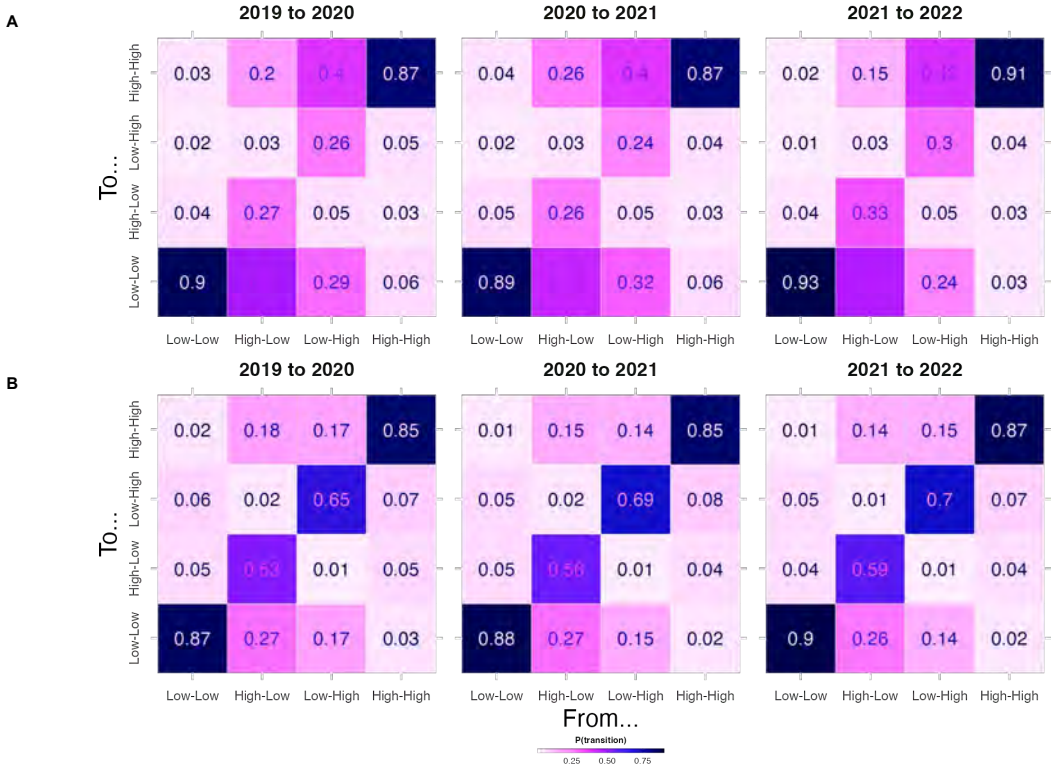

Figure A.15. **Transition probabilities.** Transition matrices for segregated clusters (**A**) and isolated clusters (**B**) showing that the hotspots of high segregation and isolation are stable over time.

We reduce dimensionality by summarizing each neighborhood using the Mann-Kendall test for monotonicity [304]. This test, which produces a statistic called Kendall's τ, captures both the direction and strength of a trend: 1 indicates perfect monotonic increase and -1 perfect monotonic decrease, with monotonicity representing the degree to which a trend is consistent, each month building on the last, compared to up-and-down over time. Fig. A.16A looks at the top 20 cities and shows that there is indeed a great deal of churn within top cities across the sample, with areas that became more isolated and more segregated in the plurality. This suggests that a rise in experienced segregation since the pandemic in 2020 is a big city phenomenon. Kendall's τ, in Fig. A.16B, shows that changes are also spatially clustered at the national level. In particular, the midwestern and northeastern parts of the country have areas that experienced steep rises in experience segregation.

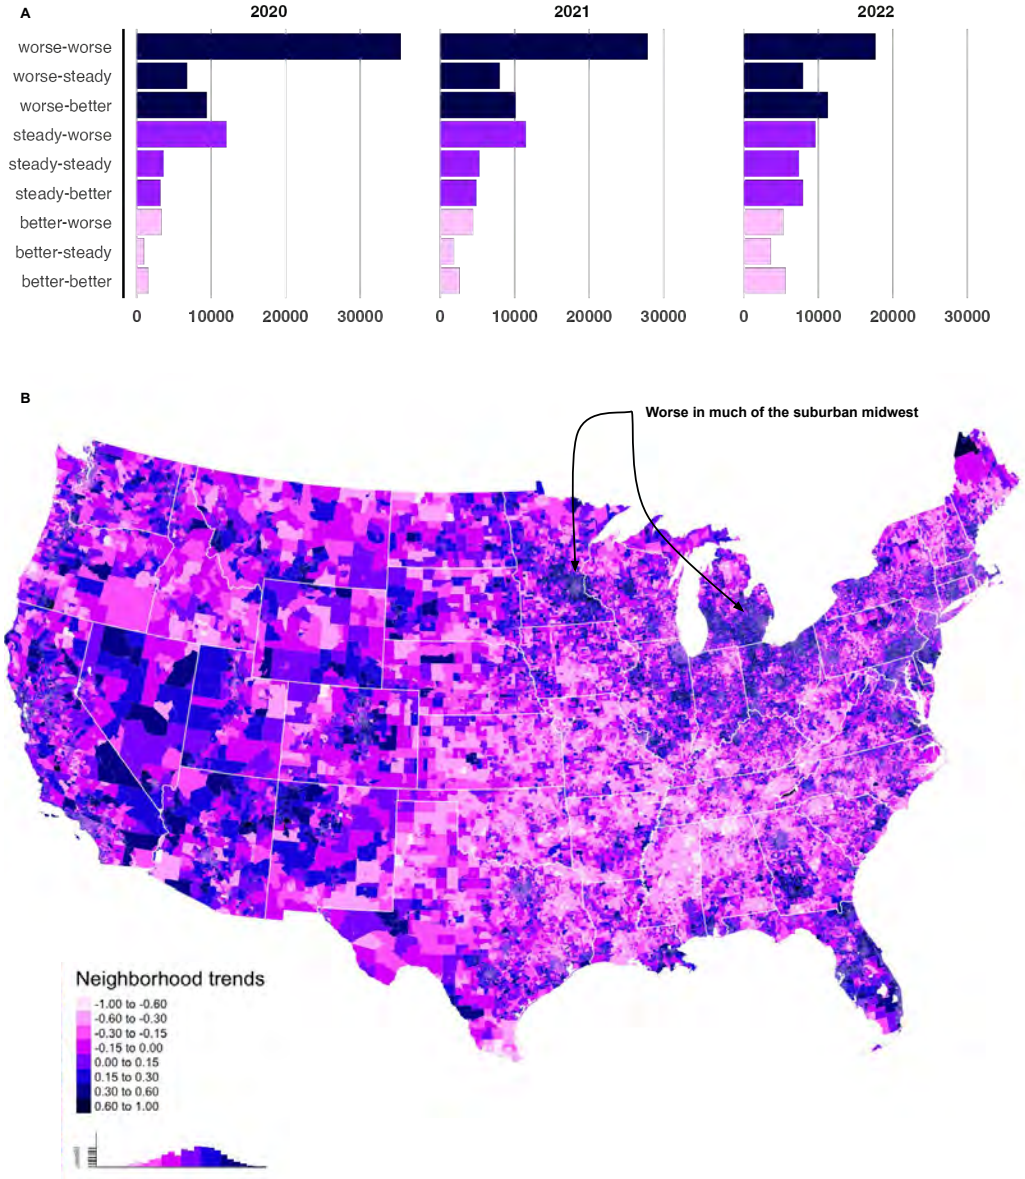

Figure A.16. **Changes over time. A** Counts of different transitions between our 3x3 classes, along segregation and isolation (first-second) for the 20 largest metros: a plurality of neighborhoods have become more segregated and isolated. **B** Trends in isolation decomposed to neighborhoods show that certain parts of the country have experienced large increases; far fewer have seen decreases in isolation.

A.15 Tree stability

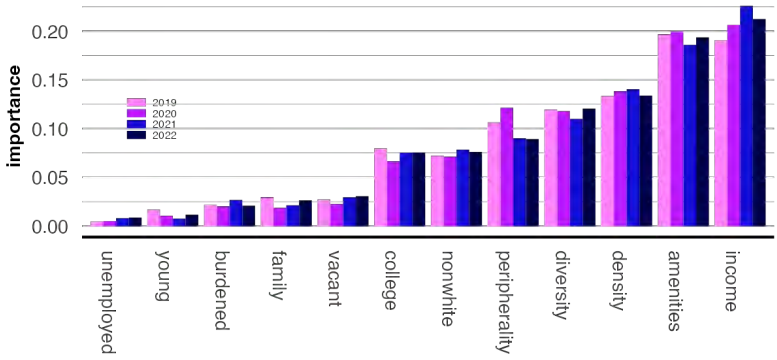

Figure A.17. **Feature importance over time.** We see remarkable stability, with the exception of a key variable: distance to the central business district. We also note that there is little change in rank importance.

A.16 Top amenities

Top Category	n
Restaurants and Other Eating Places	659,421
Offices of Physicians	486,962
Personal Care Services	385,824
Offices of Other Health Practitioners	258,376
Religious Organizations	191,314
Offices of Dentists	168,601
Automotive Repair and Maintenance	163,662
Health and Personal Care Stores	144,068
Other Amusement and Recreation Industries	139,338
Offices of Real Estate Agents and Brokers	136,412
Activities Related to Credit Intermediation	132,456
Grocery Stores	131,093
Gasoline Stations	130,596
Agencies, Brokerages, and Other Insurance Related Activities	127,865
Museums, Historical Sites, and Similar Institutions	122,852
Elementary and Secondary Schools	110,412
Accounting, Tax Preparation, Bookkeeping, and Payroll Services	101,419
Clothing Stores	95,119
Depository Credit Intermediation	93,688
Other Miscellaneous Manufacturing	87,434

Table A.7. **Top 20 Business Categories by Count.** This table presents the most common business categories in the dataset, with Restaurants and Other Eating Places being the most numerous at over 659,000 establishments. The top three categories alone account for over 1.5 million of 6 million different points of interests.

B Urban structure

B.1 Sensitivity analysis

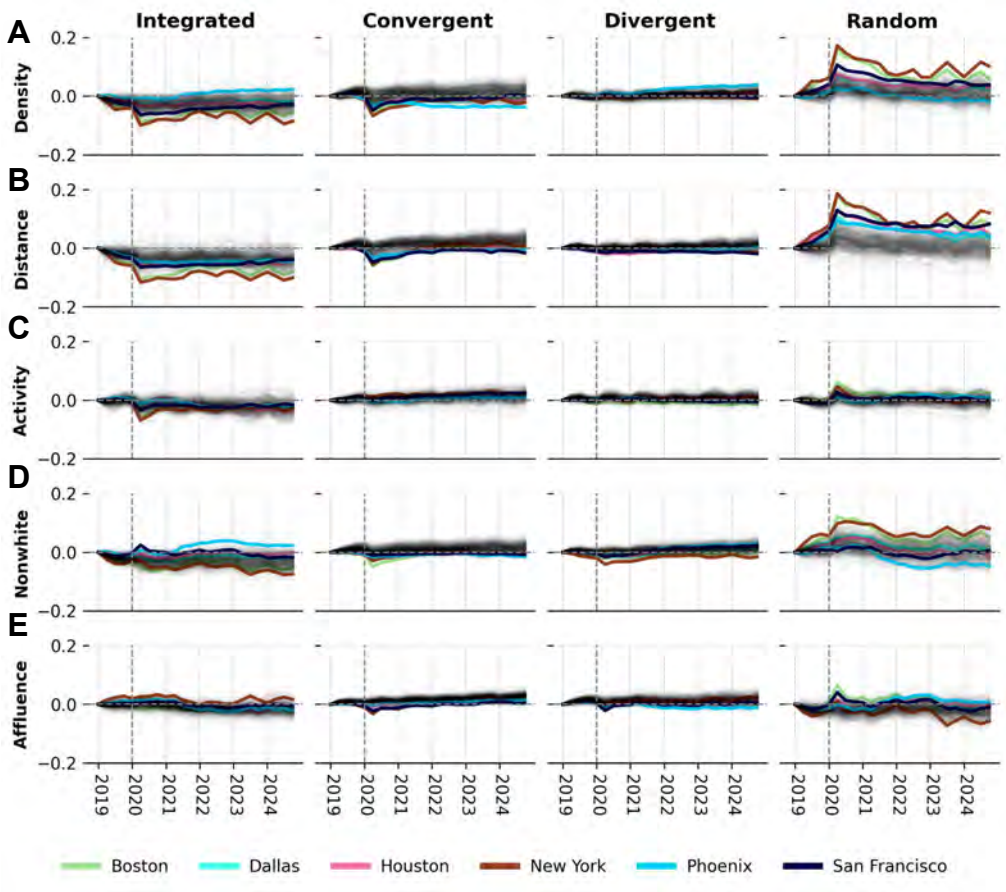

Figure B.1. **Partitioning at the median.** Splitting our matrices at the median for a given variable does little to change the results compared to the top tercile split that we use in the main text.

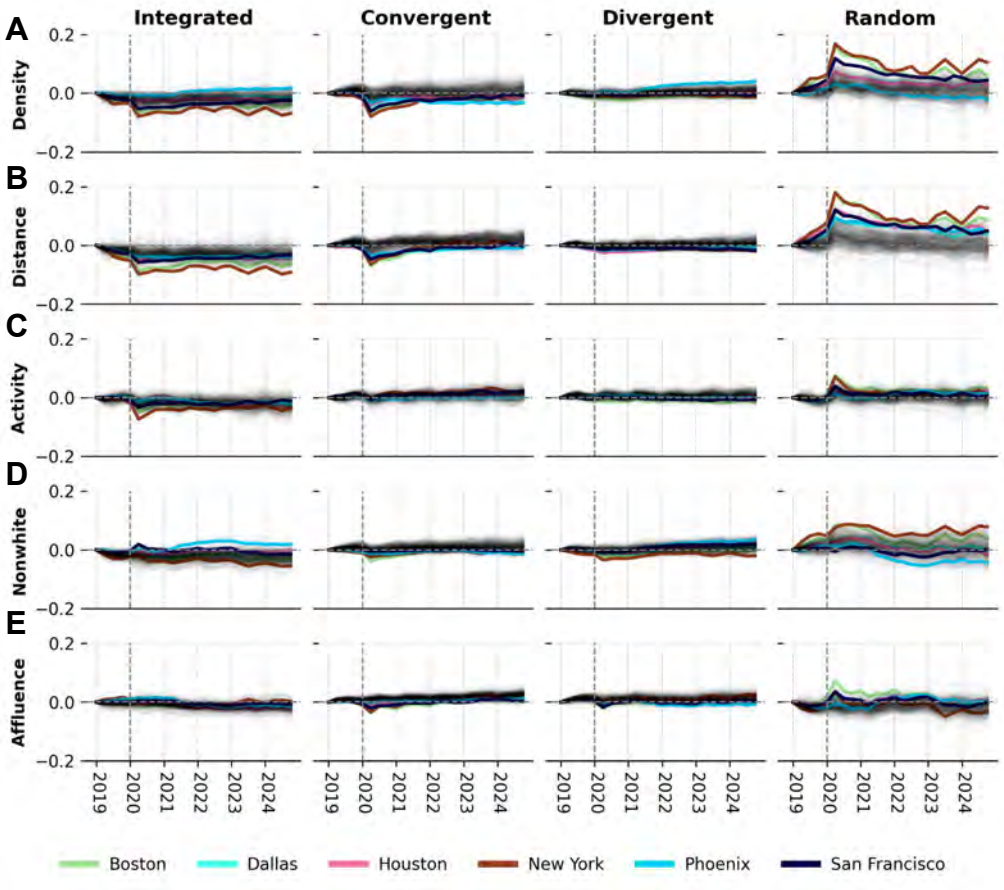

Figure B.2. **Partitioning at the top quartile.** With fewer tracts considered "hotspots" in this specification, generally the fall in integrated flows and the rise in random flows is attenuated.

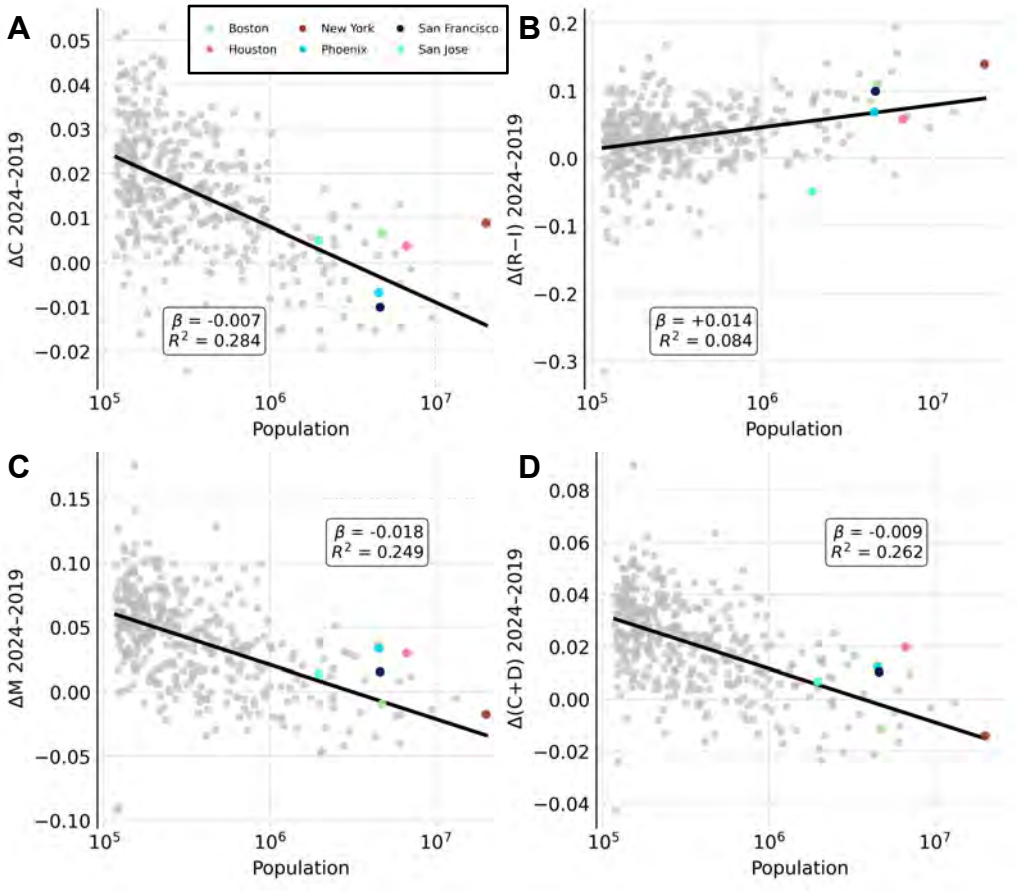

Figure B.3. **Size gradients.** We show the scatter plots the scaling relationships here. **A** There is a clear relationship between the change in convergent flows between 2019 and 2024 and the size of the city. **B** There is a shallow gradient connecting change in our doughnut index and size. **C** Change in our mixing index has a steeper gradient as does **D** the change in convergent and divergent flows in tandem.

B.2 Spending inequalities

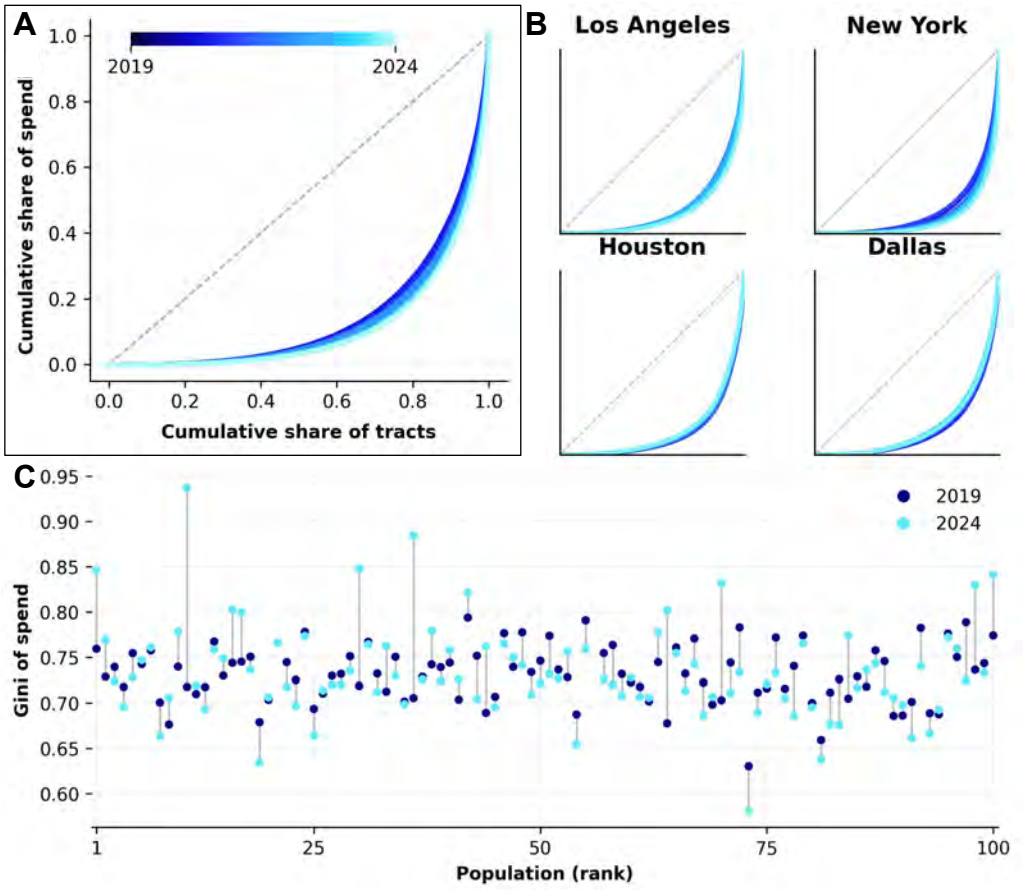

Figure B.4. **Lorenz curves and Gini coefficients. A** Pooled Lorenz curves for tract–level spend, 2019–2024 (color ramp left→right): the curve bends further from the 45° line over time, indicating rising concentration of dollars across tracts. **B** Lorenz small multiples for four metros. Inequality rises in expensive "superstar" cities (Los Angeles, New York), while Houston and Dallas show flatter or reversing curves over the same horizon. **C** Plot of Gini (spend) in *2019* against *2024* for the top–100 metros by population (rank, 1=largest). Changes are varied but exhibit no clear relationship with city size.

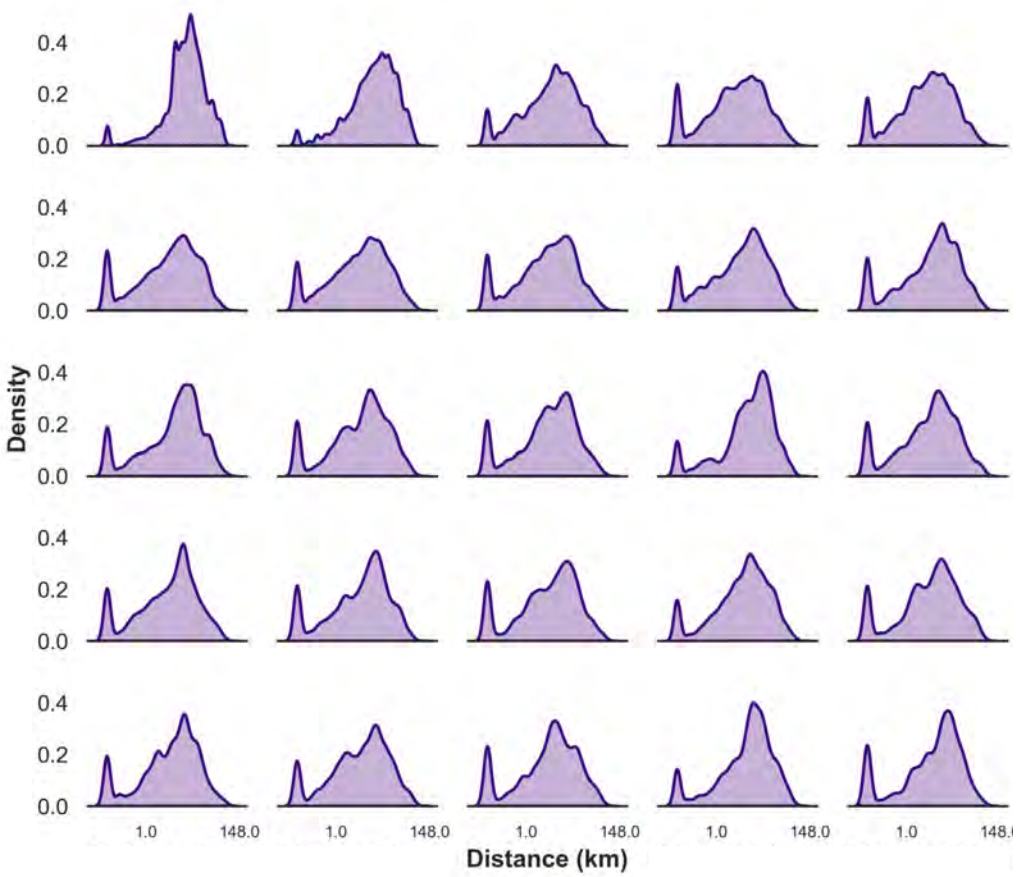

Figure B.5. **Length distributions.** Examples of distributions obtained from sampled users in our data, which show that some have obvious multimodality. Some are more bimodal, but even these distributions have spikes throughout.

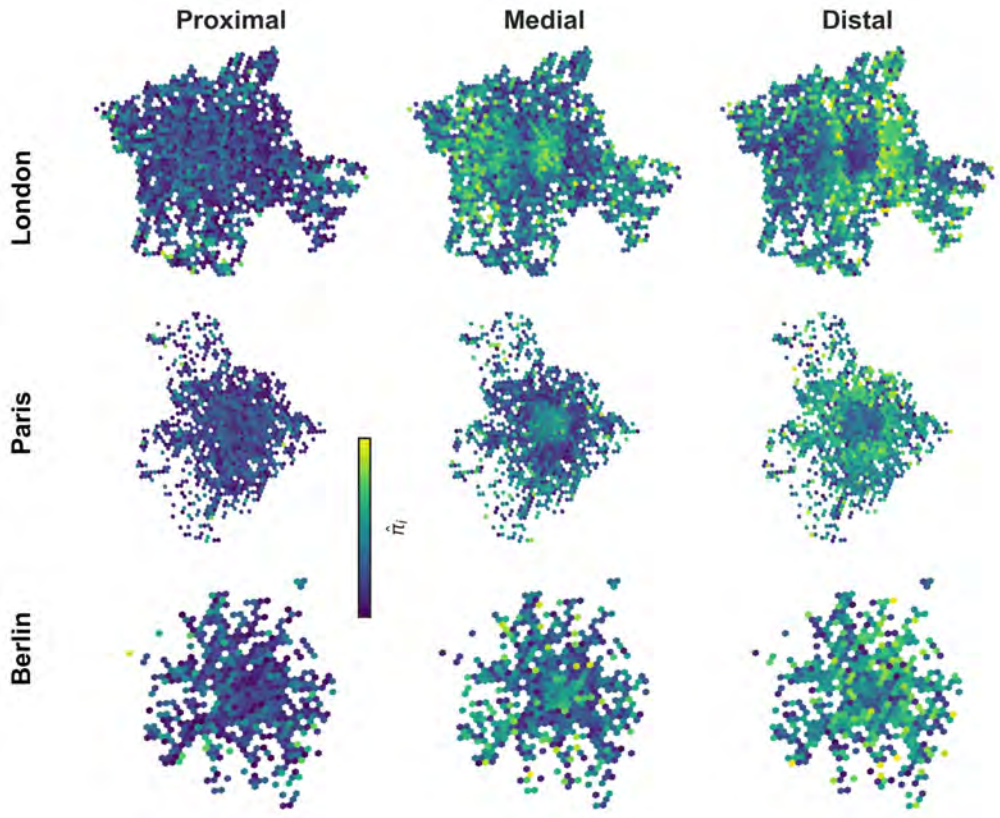

Figure B.6. **Component weights.** We note that distal weights are often lowest in the city centre, which suggests that although these areas have a distal component—because our strategy automatically assigns one—they are less important; these areas make trips that are occasional and regional, exiting the city centre.

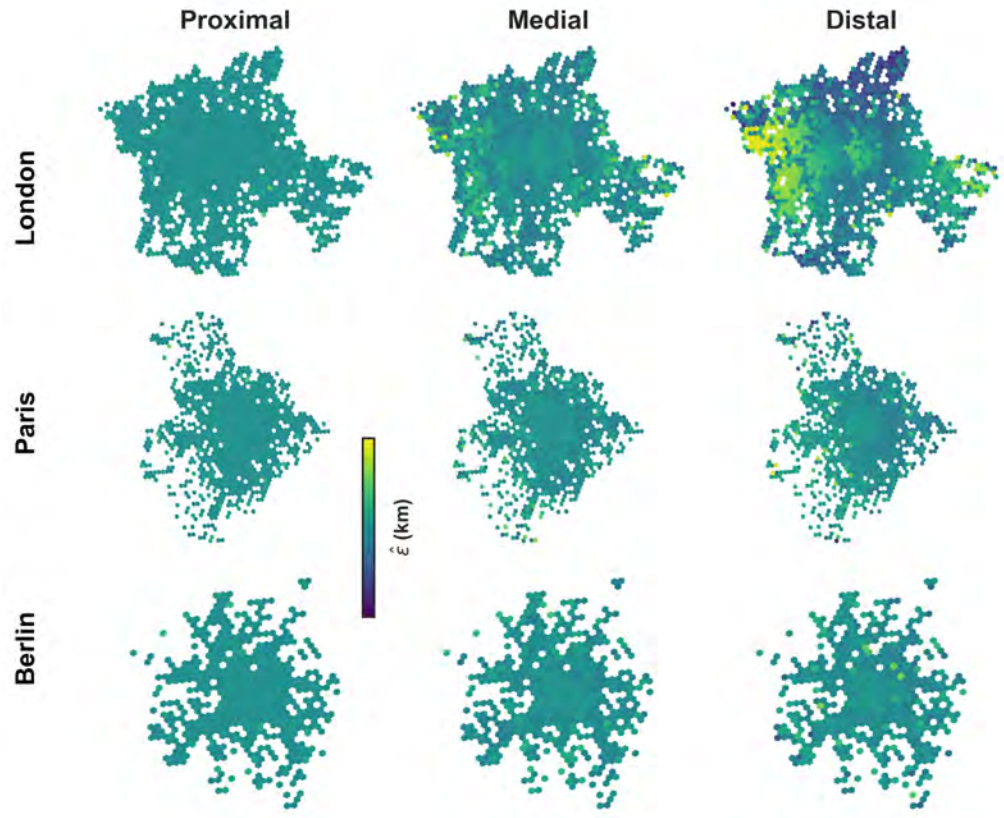

Figure B.7. **Component residuals.** Model fits are generally good across cities and components, but some areas in London have large distal errors. The most notable area here is an area that includes Slough, which is beyond the M25 and is thus a satellite of greater London.

C Heat and mobility

C.1 Data validation

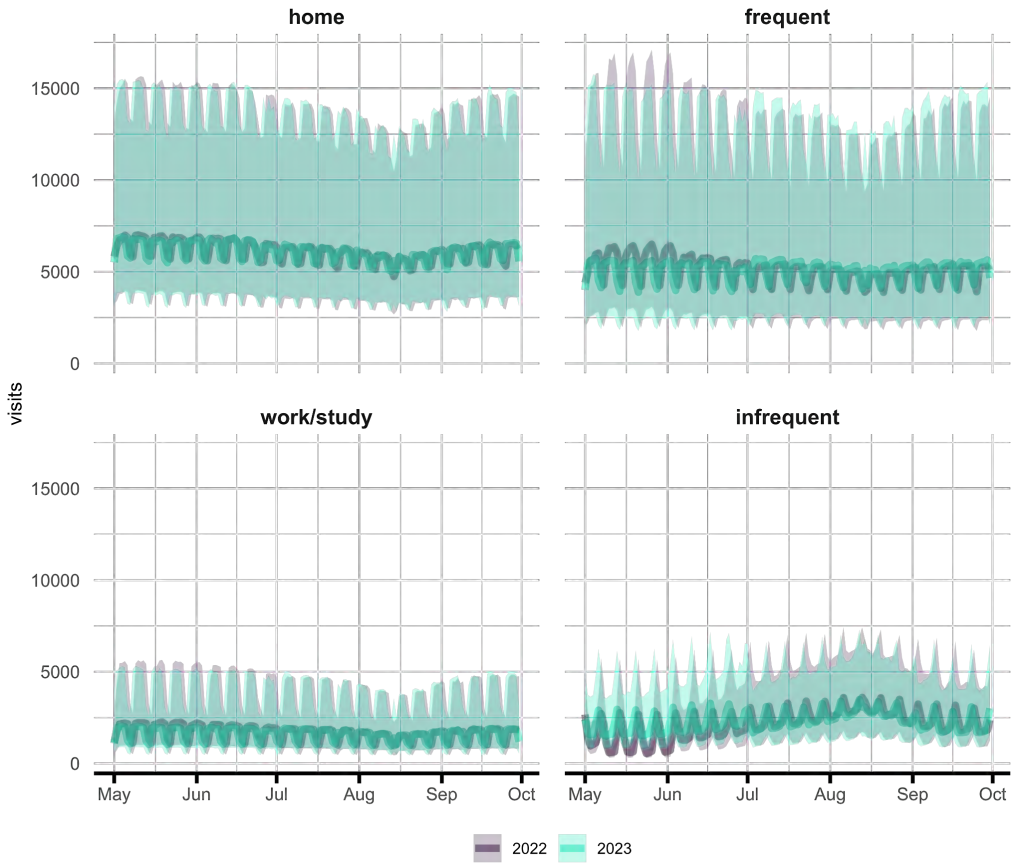

Figure C.1. **Trends in the data.** Time series per activity over the two summers of our study, with interquartile range shaded. We see that work/study and frequent activities are more common during the weak and infrequent activities are more common on the weekend; there is a seasonal trend wherein infrequent activities rise is August while work/study and frequent activities fall. We also note that in May 2022, when school is not in session, frequent activities are higher and infrequent activities lower than in the same period during 2023—possibly due to a classification error. In light of these trends, we make adjustments to our model specifications to account for these weekday/weekend and seasonal variations.

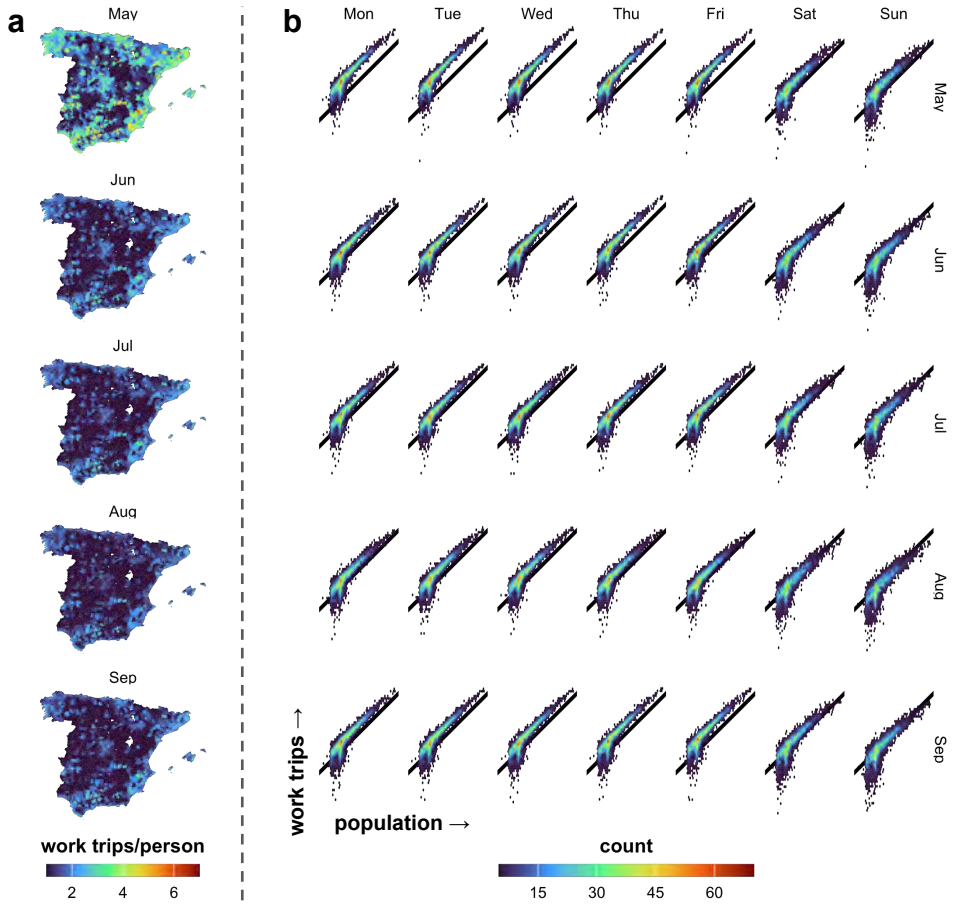

Figure C.2. **Mobility and population.** **A** The ratio work trips per person in each district for each month. We see that again in May the balance is different than in the remaining months; and urban districts tend to have more work trips than rural ones. This is because school is still in session in may, and work/study includes both of these trips. **B** We document a strong correlation between work trips and population, again with the exception of may, which typically has more work trips per person than other months. The fit returns to that of other months on the weekends, when school is out, during May. We address this in the model specification, and we also perform sensitivity analysis including and excluding this period.

C.2 Model results

Table C.1. Impact of Extreme Heat (UTCI $> 40°$C) on Mobility by Activity Type

Dependent Variable:	*Flows*	
	Within districts	Between districts
Variables		
Work or Study	-0.0068	-0.0054
	(0.0087)	(0.0079)
Frequent Activity	-0.0207**	-0.0243***
	(0.0086)	(0.0069)
Infrequent Activity	-0.125***	-0.105***
	(0.0311)	(0.0232)
Fixed effects		
District	Yes	Yes
Date	Yes	Yes
Fit statistics		
Pseudo R^2	0.978	0.972

Notes: Standard errors (in parentheses) are clustered by province.
Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Note: Results from Poisson regressions including district and date fixed effects. Estimates represent semi-elasticities of flows (i.e., logs of expected counts). Standard errors are robust to clustering at the province level.

C.3 Sensitivity analysis

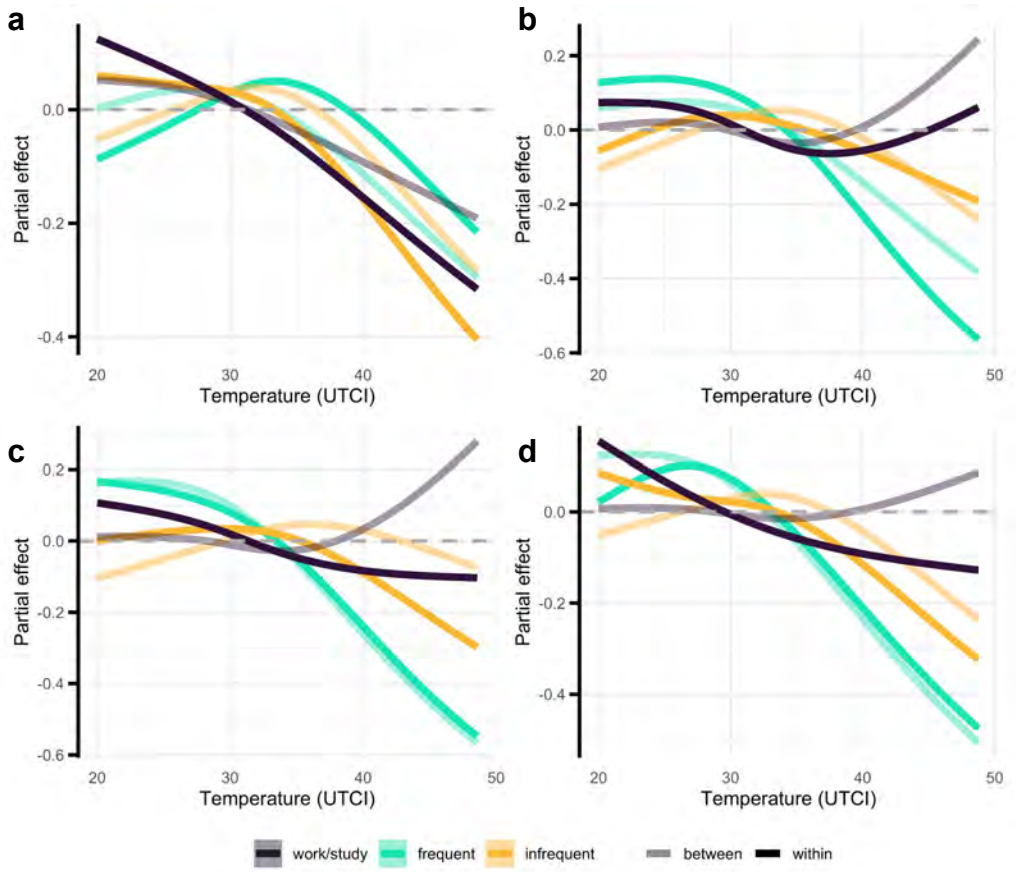

Figure C.3. **Sensitivity analysis.** **A** GAM results for the top 50 cities in Spain, all with a population over 100,000, and **B** for the remainder of Spanish cities, showing that the results are largely stable, with the possible exception being around work patterns: in large cities, people are more likely to avoid travel to work on hot days than in small cities—interesting because information technology jobs concentrate in large cities and thus can likely be taken from home [182]. **C** and **D** filter the data to 2022 and 2023, respectively, showing that they are also stable across time.

C.4 Placebo tests

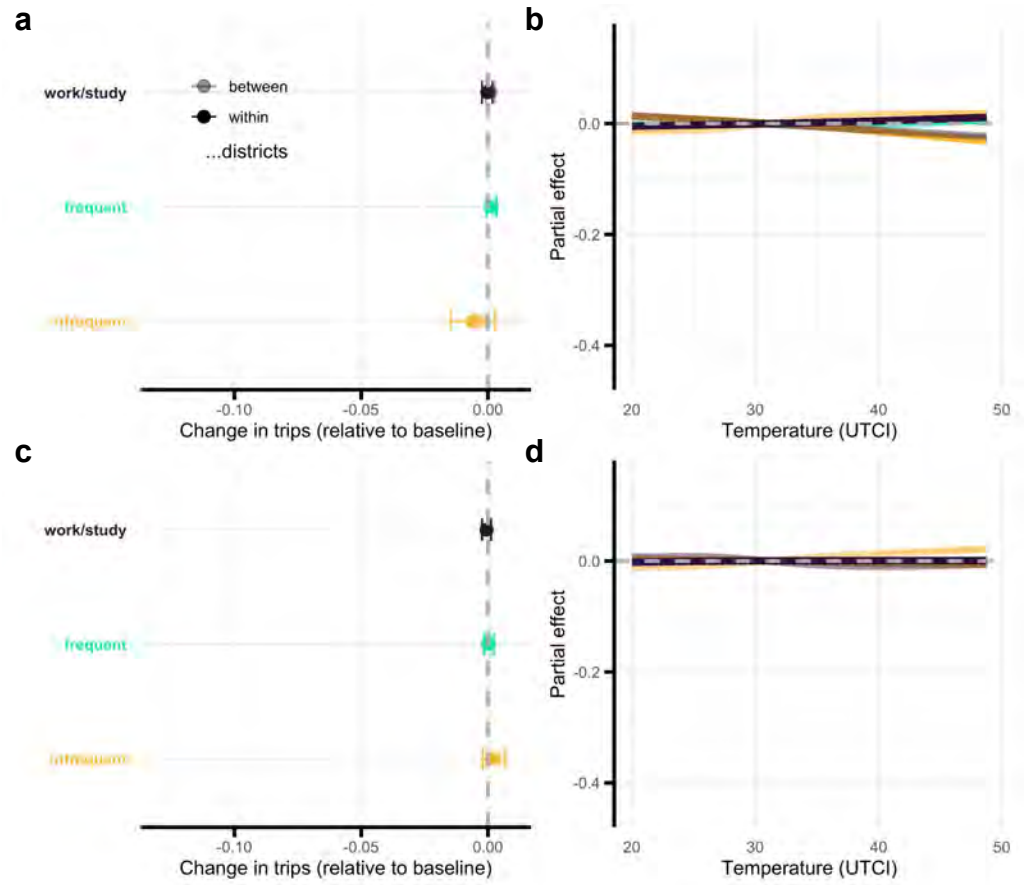

Figure C.4. **Placebo tests.** We perform 4 separate placebo tests to rule out potential spurious associations: **A** and **B** shuffle temperatures across districts on the same date, while **C** and **D** shuffle temperatures across different dates within the same district. The cross-district shuffles help rule out the possibility that our results are driven by events affecting all districts simultaneously. The within-district shuffles help rule out the possibility that observed changes would occur within districts regardless of temperature variation. While GAM confidence intervals occasionally exclude zero, effect sizes are reduced by an order of magnitude or more compared to the main analysis, supporting our primary findings.

C.5 Consistency between genders

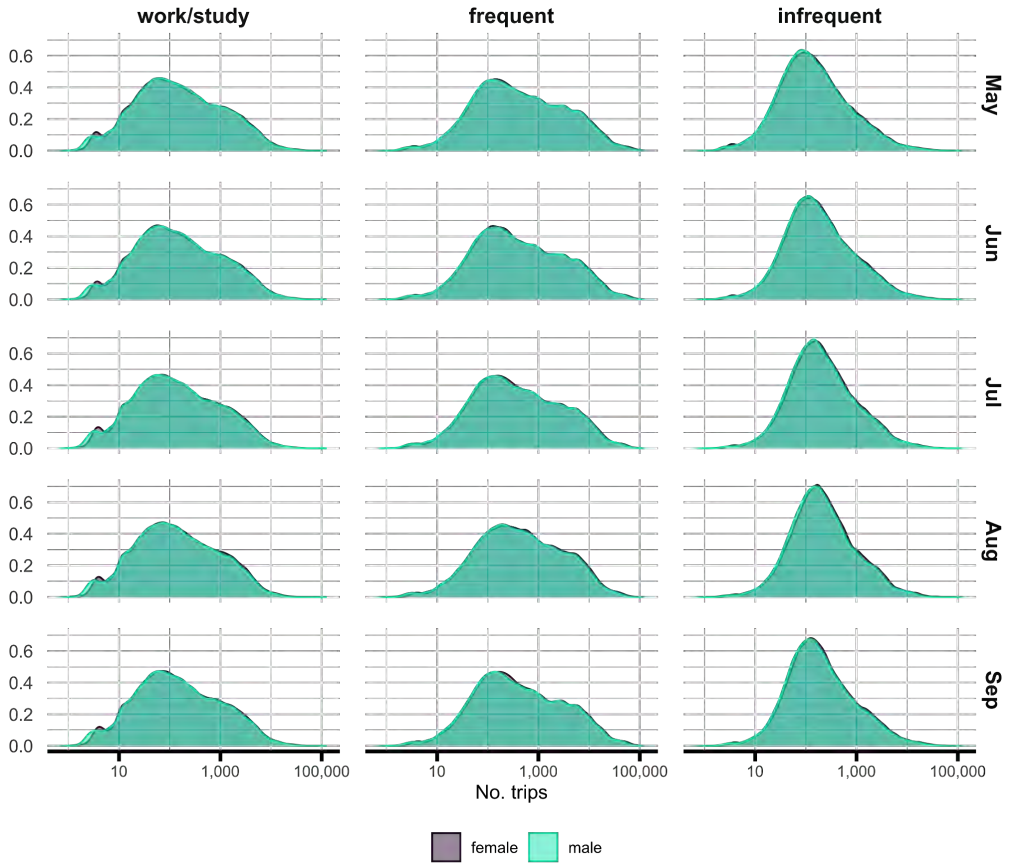

Figure C.5. **Trip distributions.** Here we show the distributions of daily visits per month and class of trip. We note that across all months and classes of trip, the distributions of trips are consistent between genders, suggesting broad similarity.

C.6 Gravity model

Dependent Variable:	Trips		
	Base model	Temperature bins	With gradient
Model:	(1)	(2)	(3)
Variables			
Log Distance	-1.87***	-1.87***	-1.87***
	(0.013)	(0.013)	(0.013)
Temperature 25-30°C		0.011***	0.012***
		(0.002)	(0.002)
Temperature 30-35°C		0.008**	0.009***
		(0.004)	(0.004)
Temperature 35-40°C		-0.020***	-0.018***
		(0.006)	(0.006)
Temperature 40-45°C		-0.053***	-0.051***
		(0.007)	(0.007)
Temperature 45-50°C		-0.066***	-0.062***
		(0.009)	(0.010)
Log Distance × Temperature Gradient			-0.001*
			(0.0007)
Fixed-effects			
ID origin	Yes	Yes	Yes
ID destination	Yes	Yes	Yes
date	Yes	Yes	Yes
Fit statistics			
Observations	31,148,196	31,148,196	31,148,196
Pseudo R^2	0.87382	0.87386	0.87386
Squared Correlation	0.78625	0.78636	0.78637

Clustered (origin & destination) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table C.2. **Impact of temperature on mobility.** We use gravity models to understand if flows are displaced, moving from one district to another, during heat waves, which would violate TWFE assumptions. While high temperatures reduce flows, agreeing with the TWFE analysis, we see that temperature gradient has no significant effect, suggesting that people do not preference higher or lower temperatures during heat waves. This also indicates that people do not seek out cooler parts of the country during extreme heat. All models include origin, destination, and date fixed effects; Standard errors clustered by origin and destination.

	Population				
Temperature	Very Low	Low	Medium	High	Very High
25-30°C	0.0267***	0.0324***	0.0321***	0.0376***	0.0482***
	(0.0031)	(0.0045)	(0.0086)	(0.0091)	(0.0092)
30-35°C	0.0589***	0.0629***	0.0371***	0.0277**	0.0224**
	(0.0049)	(0.0059)	(0.0078)	(0.0080)	(0.0084)
35-40°C	0.0862***	0.0834***	0.0350***	-0.0007	-0.0363***
	(0.0075)	(0.0095)	(0.0077)	(0.0076)	(0.0077)
40-45°C	0.0791***	0.0763***	0.0086	-0.0303**	-0.0874***
	(0.0081)	(0.0108)	(0.0083)	(0.0089)	(0.0087)
45-50°C	0.0751***	0.0588***	0.0123	-0.0783***	-0.0962***
	(0.0158)	(0.0125)	(0.0130)	(0.0173)	(0.0146)
log(d)	-1.865*** (0.0112)				
Fixed-effects					
ID origin	Yes				
ID destination	Yes				
date	Yes				
Fit statistics					
Observations	19,638,009				
Pseudo R^2	0.87786				
Squared Correlation	0.79586				

Clustered (by origin & destination) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, .: 0.1*

Table C.3. **Impact of temperature and combined population on mobility.** In this gravity model we report the results of a gravity model that groups flows by population at origin and destination. Because the model has 5x5 population classes as well as temperature bands, we present the average effects for combined population groups. Our results show large negative effects for the most populous areas, while moderate warmth is positive relative to cold.

C.7 Projecting into the future

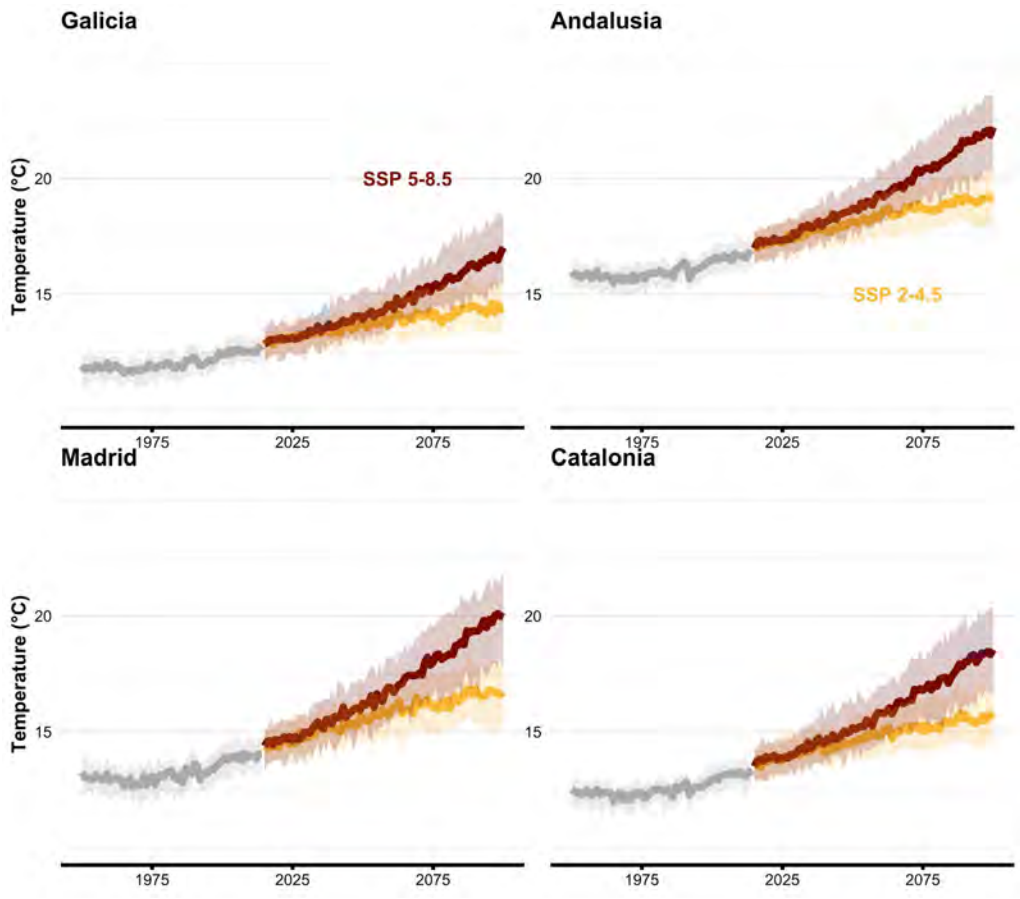

Figure C.6. **CMIP estimates.** Andalusia, in the South will see annual temperatures rise more than Galicia, in the North. Major cities like Barcelona and Madrid will also experience considerable changes.

C.8 Network properties

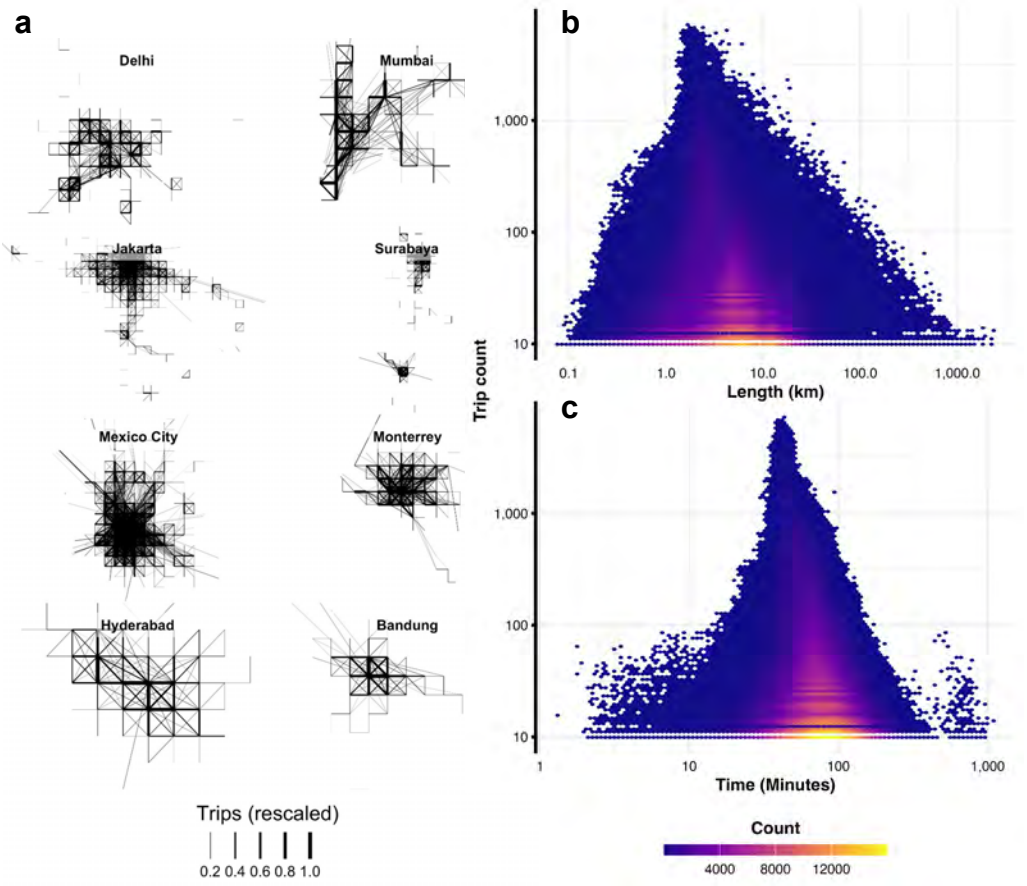

Figure C.7. **Networks and statistics. A** We construct networks where edges correspond to trips and nodes corresponds to geohash5 aggregations. **B** and **C** show the properties of these edges, with trip times averaging between 50 and 100 minutes and lengths in kilometers averaging between 5 and 10. We note a distance and time decay, where trips of long lengths and times are less frequent across the sample.

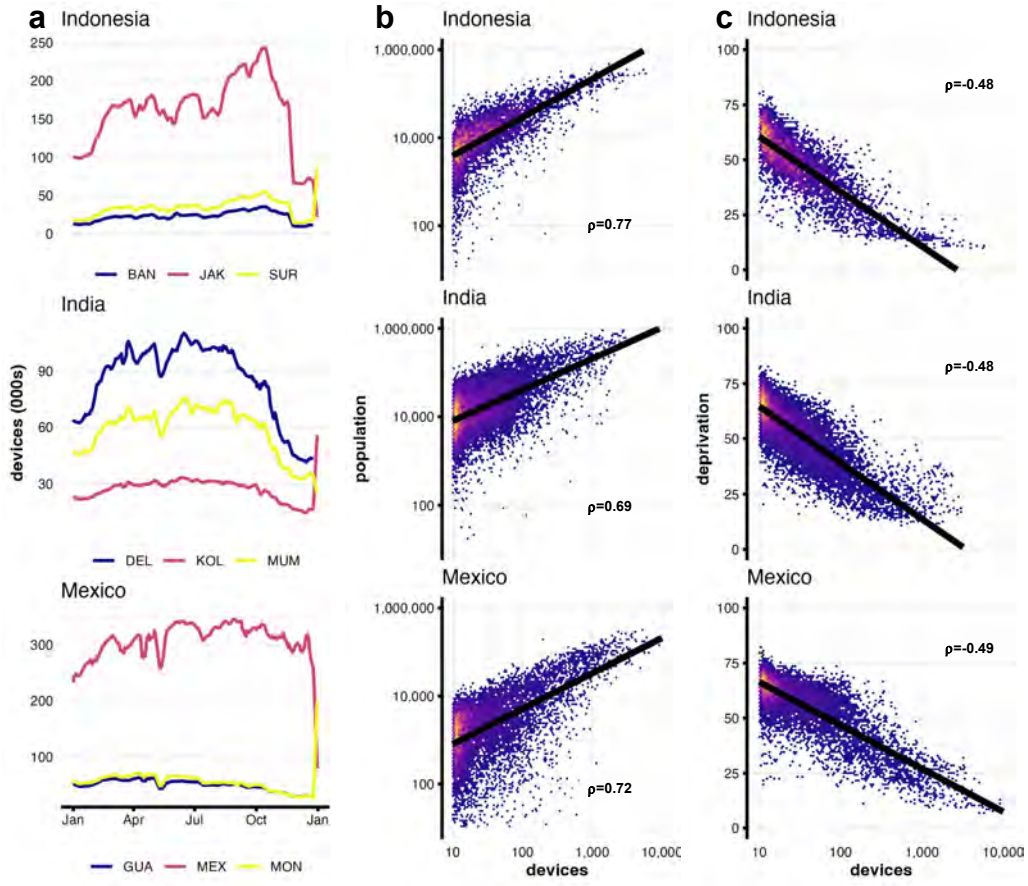

Figure C.8. **Biases in the data.** **A** We show the time series for various top cities in the data, finding irregularities in the beginning and end of the year, which we remove from the data. We also follow the recommendation the data provider [533] and remove a number of other anomalous days in May and August. **B** and **C** show how device counts in each geohash correlate with population and deprivation. While see a strong correlation to population by geohash in our data, lending confidence to our estimates, we also see that more deprived areas are less represented in the data. Because population itself is anticorrelated with deprivation ($\rho = -0.7$), this is not entirely the product of bias, but we do our best to address this with stratification and sensitivity analysis.

Origin-Destination Patterns	Trips	Share (%)
India		
Urban-rural	4,527	0.01%
Rural-urban	4,870	0.01%
Inter-urban	30,152	0.05%
Rural-rural	2,333,024	3.54%
Within-urban	63,527,103	96.40%
of which within-cell	52,383,732	79.49%
Total trips	65,899,676	100.00%
Indonesia		
Urban-rural	6,952	0.02%
Rural-urban	8,234	0.02%
Inter-urban	16,182	0.04%
Rural-rural	1,265,869	3.04%
Within-urban	40,366,663	96.89%
of which within-cell	30,211,502	72.51%
Total trips	41,663,900	100.00%
Mexico		
Urban-rural	221,369	0.17%
Rural-urban	226,830	0.18%
Inter-urban	129,046	0.10%
Rural-rural	15,516,385	12.06%
Within-urban	112,561,977	87.49%
of which within-cell	76,076,308	59.13%
Total trips	128,655,607	100.00%

Table C.4. **Comparison of trip patterns across all three countries.** Most trips in all countries are urban-urban and very few begin or end in rural areas. In Mexico, the most developed country in our dataset, the share of trips originating and terminating the same cell is the lowest, indicating greater mobility.

C.9 Alternative specifications

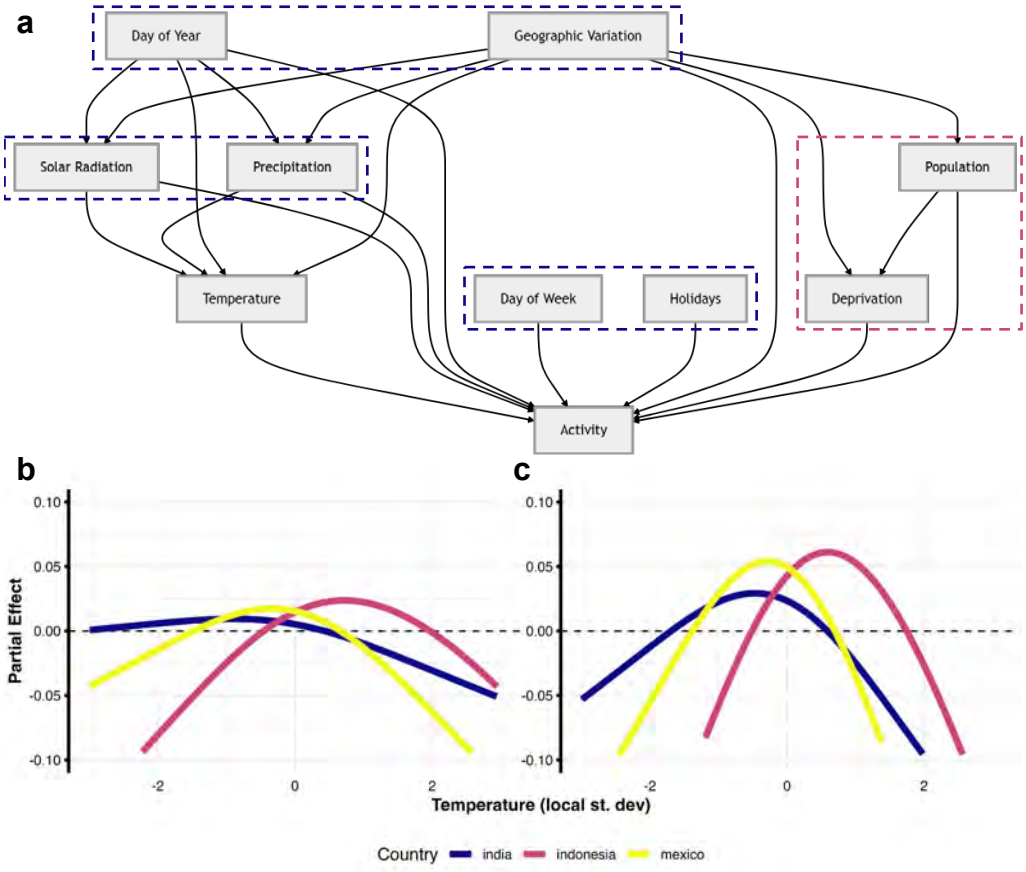

Figure C.9. **Modeling temperature and activity. A** We set up a directed acyclic graph (DAG) to ensure that we close all necessary causal paths to activity; we control for day-of-year, day-of-week, holidays, solar radiation, precipitation, and use geographic fixed effects that necessarily stratify by population and deprivation. **B** Model results for different countries, showing that the highest activity levels occur at average temperatures for an area, and that high extremes correspond with fewer trips. **C** Model results using trip duration rather than trip count reveal that the longest trips tend to occur at average temperatures, with extreme high temperatures leading to shorter trips.

C.10 Placebo tests

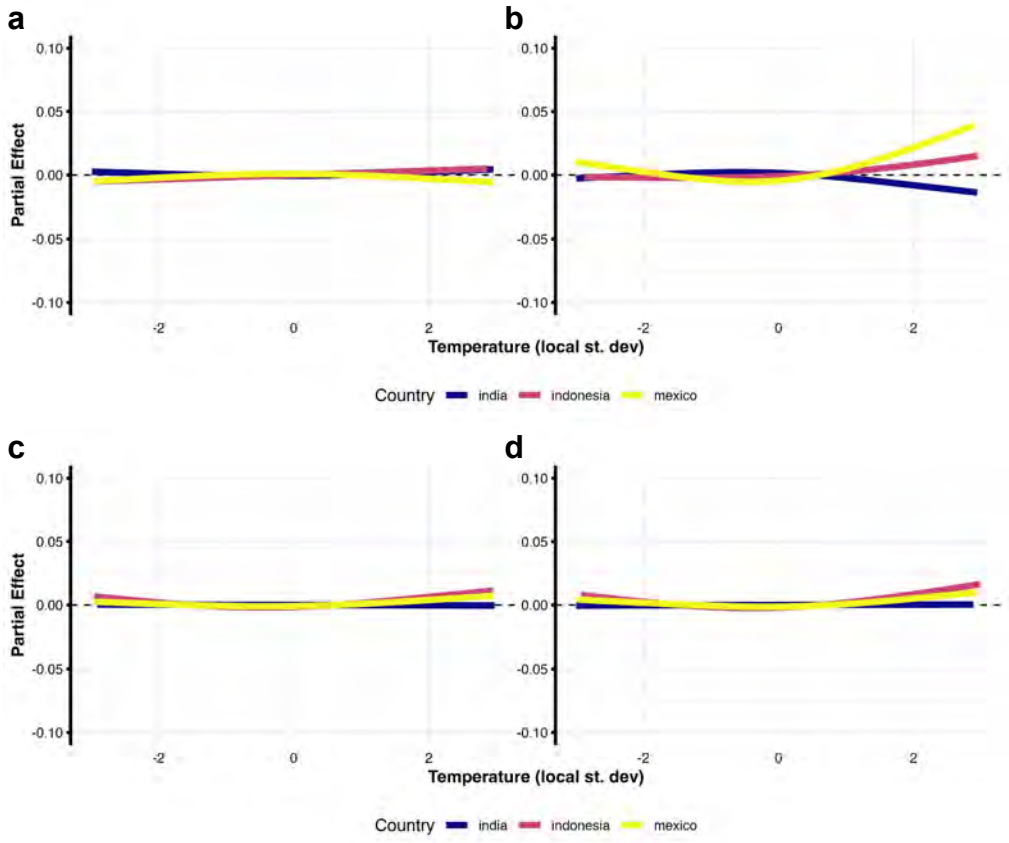

Figure C.10. **Permutations tests.** We use a variety of permutation tests to ensure that the relationship we observed is not spurious, including shuffled temperatures within GeoHash5 between dates, giving each cell a temperature from a different time of year in that cell. **A** and **B** does this for duration and trip count, respectively. We also shuffled between GeoHash5 within dates, giving each cell a new temperature from the same date in a different cell, again for duration and trip count in **C** and **D**.

C.11 Sensitivity analysis

Constraint Level and Requirements	No. of cells	No. of cities
India		
Low constraint: Min. days with mobility=50 Min. average mobility=10	1,259	482
Medium constraint: Min. days with mobility=100 Min. average mobility=15	1,043	384
High constraint: Min. days with mobility=150 Min. average mobility=20	832	262
Indonesia		
Low constraint: Min. days with mobility=50 Min. average mobility=10	701	199
Medium constraint: Min. days with mobility=100 Min. average mobility=15	574	161
High constraint: Min. days with mobility=150 Min. average mobility=20	461	103
Mexico		
Low constraint: Min. days with mobility=50 Min. average mobility=10	964	177
Medium constraint: Min. days with mobility=100 Min. average mobility=15	890	176
High constraint: Min. days with mobility=150 Min. average mobility=20	731	109

Table C.5. **Sample sizes under varying data quality constraints.** We test our results on various samples of the data, conditioning on the level of mobility in the cells to ensure that large changes on small values are not driving our results.

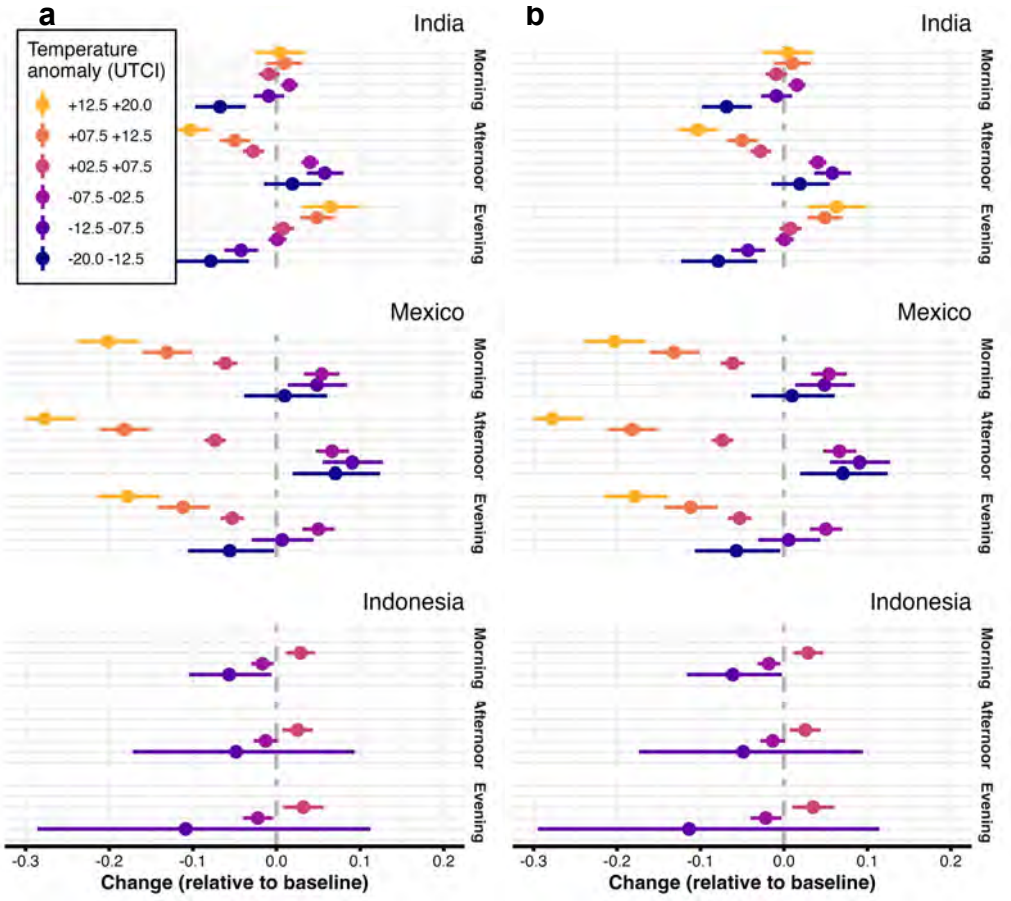

Figure C.11. **Medium and high constraints.** We show the results of our TWFE regressions for both **A** medium and **B** high constraints respectively. These results, which are almost identical, show that different subsets of the data behave in similar ways, and our results are not driven by certain groups of cells with high rates of change on low values.

C.12 Time-of-day and population

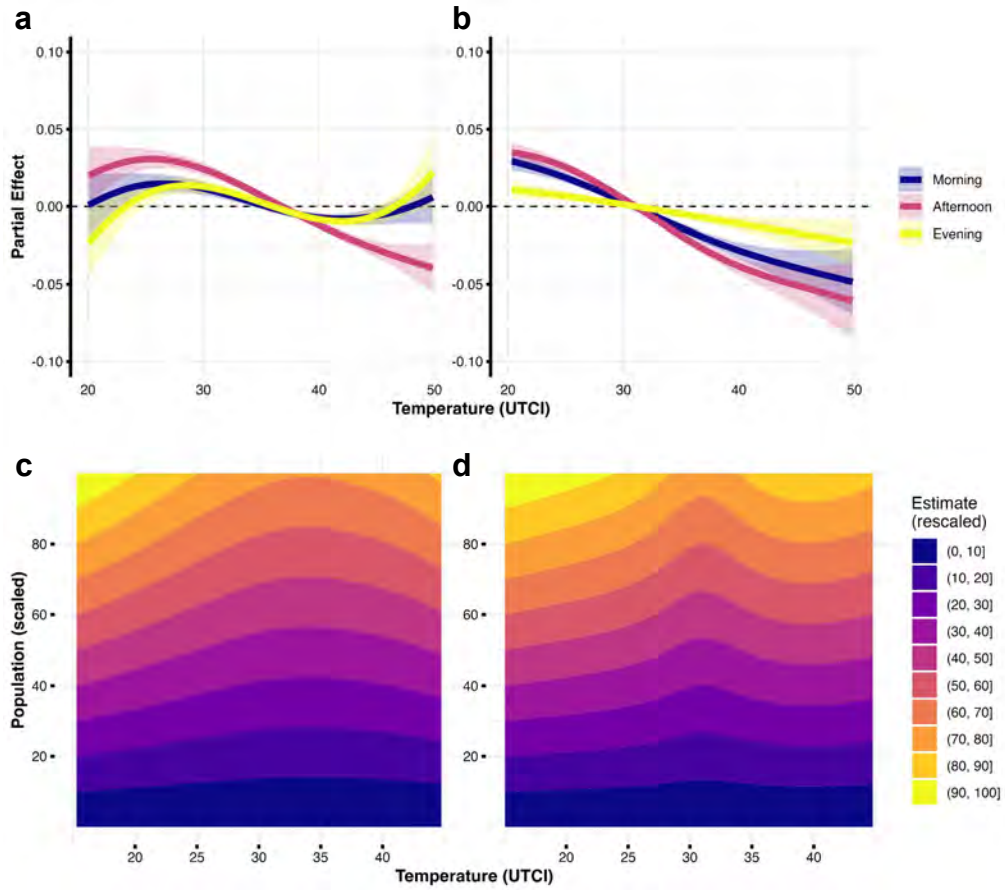

Figure C.12. **Mobility by time of day and population.** **A** When we decompose the effect of temperature on activity by time of day, we see that the reduction in activity comes during the afternoons, suggesting that people are avoiding the hottest part of the day. **B** In Mexico, all periods of the day see reductions in activity. We see that in India **C** and Mexico **D**, larger populations see stronger temperature effects in absolute terms, with higher predicted values and moderate temperatures and lower predicted values at the extremes, while these curves are flattened for smaller populations.

C.13 ARIMAX curves

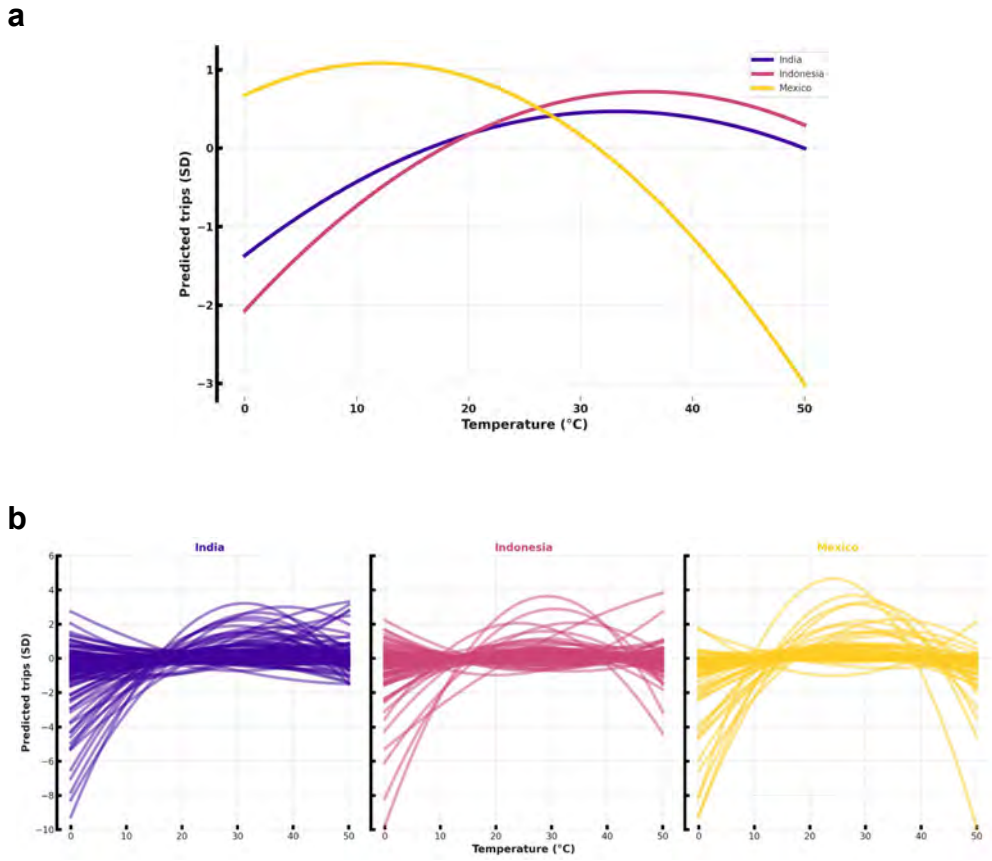

Figure C.13. **ARIMAX curves. A** We plot the curves for optimal temperature given the coefficients from the average fitted ARIMAX model, seeing that the predictions turn negative between 30 and 40°C for both India and Indonesia, but fall much earlier in Mexico—which has a temperate climate many populous areas like Mexico City and Guadalajara. **B** We disaggregate those averages to show all curves from all geohash3 ARIMAX models: there are some curves that show a positive effect of temperature but most show a negative one, in particular those fit in India.

C.14 Projections for Indonesia

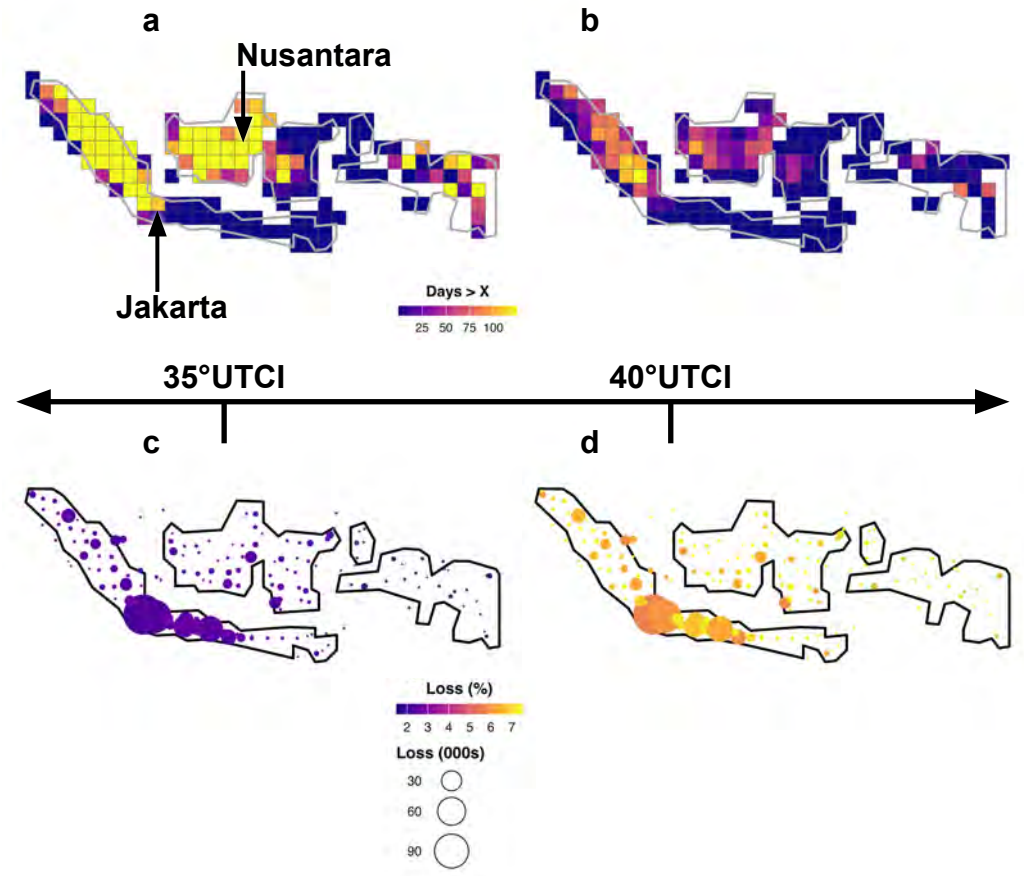

Figure C.14. **Temperatures and activity in 2050. A** and **B** show frequency of days above $35°$UTCI and $40°$UTCI, respectively; we see that Borneo and Sumatra will have more frequent days above $35°$UTCI while days $40°$UTCI will be concentrated in Sumatra, suggesting that heat ways will be a problem there. We note here that Indonesia is in the process of moving its capital city to a hotter part of the country, which will either force adaptation for new residents or limit activity. **C** and **D** show the consequences of these temperatures according to our model, with strong effects at higher temperatures but Indonesia, the existing capital and largest city, is will show the largest effects in absolute terms.

C.15 CMIP6 estimates under different scenarios

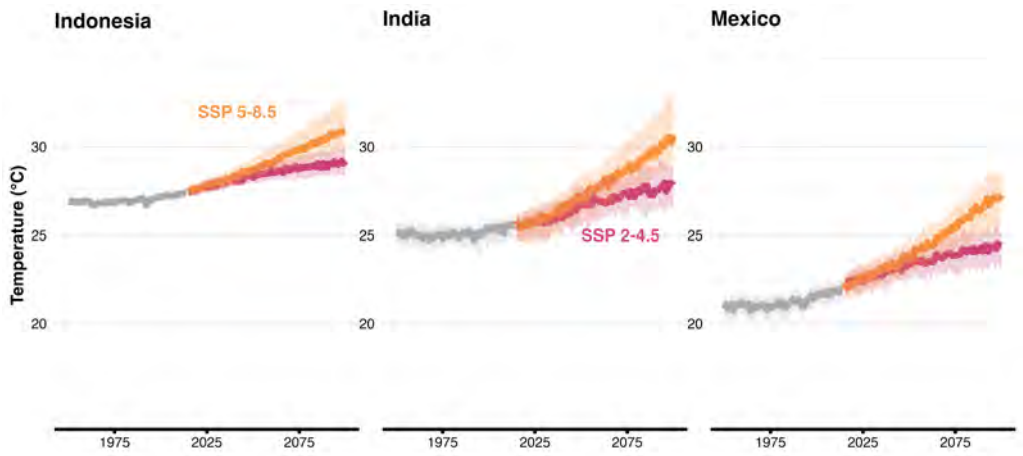

Figure C.15. **Temperatures under SSP 4.5 and 8.5 through 2100.** We show CMIP6 estimates under different scenarios for each country, with Mexico seeing warming from a low average today, India seeing the strongest change, and Indonesia seeing less of a change but from a high average today. These estimates are the averages from ensembles of models and the ribbons indicate the 10th and 90th percentiles of all models [515].

References

[1] J. T. Abatzoglou and A. P. Williams. Impact of anthropogenic climate change on wildfire across western us forests. *PNAS*, 113(42):11770–11775, 2016.

[2] S. Abbasi, J. Ko, and J. Min. Measuring destination-based segregation through mobility patterns: Application of transport card data. *Journal of Transport Geography*, 92:103025, 2021.

[3] T. Abbiasov, C. Heine, E. L. Glaeser, C. Ratti, S. Sabouri, A. S. Miranda, and P. Santi. The 15-minute city quantified using mobility data. Technical report, National Bureau of Economic Research, 2022.

[4] M. G. Abdelmonem and R. McWhinney. In search of common grounds: Stitching the divided landscape of urban parks in belfast. *Cities*, 44:40–49, 2015.

[5] A. Aberra, G. A. Aga, F. Jolevski, and N. Karalashvili. Understanding informality: Comprehensive business-level data and descriptive findings. Technical report, The World Bank, 2022.

[6] Advan. "products and patterns". `https://advanresearch.com/products-patterns`, 2024.

[7] P. Aghion and P. Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992.

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

[9] L. M. Aiello, A. Vybornova, S. Juhász, M. Szell, and E. Bokányi. Urban highways are barriers to social ties. *Proceedings of the National Academy of Sciences*, 122(10):e2408937122, 2025.

[10] Airports Council International (ACI) World. The busiest airports in the world defy global uncertainty and hold top rankings, April 2025. Press release. States that the Top 10 busiest airports represented 9̃% of global passenger traffic in 2024; Top 10 aircraft movements 6̃%.

[11] Ajuntament de Barcelona. Refugis climàtics: mapa i xarxa municipal, 2025.

[12] C. G. Aksoy, J. M. Barrero, N. Bloom, S. J. Davis, M. Dolls, and P. Zarate. Working from home around the world. *Brookings Papers on Economic Activity*, 53(2):281–360, 2022. Data from the Global Survey of Working Arrangements (G-SWA).

[13] C. G. Aksoy, J. M. Barrero, N. Bloom, S. J. Davis, M. Dolls, and P. Zarate. Working from home around the world. *Brookings Papers on Economic Activity*, 2022(2):281–360, 2022.

[14] C. G. Aksoy, J. M. Barrero, N. Bloom, S. J. Davis, M. Dolls, and P. Zarate. Time savings when working from home. Technical Report 30866, National Bureau of Economic Research, 2023.

[15] C. G. Aksoy, J. M. Barrero, N. Bloom, S. J. Davis, M. Dolls, and P. Zarate. The global persistence of work from home. *Proceedings of the National Academy of Sciences*, 122(27):e2509892122, 2025.

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

[17] L. Alessandretti, U. Aslak, and S. Lehmann. The scales of human mobility. *Nature*, 587(7834):402–407, Nov. 2020. Number: 7834 Publisher: Nature Publishing Group.

[18] L. Alessandretti, P. Sapiezynski, V. Sekara, S. Lehmann, and A. Baronchelli. Evidence for a conserved quantity in human mobility. *Nature human behaviour*, 2(7):485–491, 2018.

[19] J.-V. Alipour, O. Falck, S. Krause, C. Krolage, and S. Wichert. Working from home and

consumption in cities. Technical report, CESifo Working Paper, 2025.
[20] G. W. Allport and B. M. Kramer. Some roots of prejudice. *The Journal of Psychology*, 22(1):9–39, 1946.
[21] M. Almagro, J. Coven, A. Gupta, and A. Orane-Hutchinson. Racial Disparities in Frontline Workers and Housing Crowding during COVID-19: Evidence from Geolocation Data. Opportunity and Inclusive Growth Institute Working Papers 37, Federal Reserve Bank of Minneapolis, Sept. 2020.
[22] W. Alonso. *Location and Land Use*. Harvard University Press, 1964.
[23] E. G. Altmann. Spatial interactions in urban scaling laws. *Plos one*, 15(12):e0243390, 2020.
[24] L. G. A. Alves, H. V. Ribeiro, E. K. Lenzi, and R. S. Mendes. Distance to the Scaling Law: A Useful Approach for Unveiling Relationships between Crime and Urban Metrics. *PLOS ONE*, 8(8):e69580, Aug. 2013. Publisher: Public Library of Science.
[25] E. O. Ananat. The wrong side (s) of the tracks: The causal effects of racial segregation on urban poverty and inequality. *American Economic Journal: Applied Economics*, 3(2):34–66, 2011.
[26] A. Anas, R. Arnott, and K. A. Small. Urban spatial structure. *Journal of Economic Literature*, 36(3):1426–1464, 1998.
[27] F. Andersson, S. Burgess, and J. I. Lane. Cities, matching and the productivity gains of agglomeration. *Journal of Urban Economics*, 61(1):112–128, 2007.
[28] A. L. Angert. The space-for-time gambit fails a robust test. *Proceedings of the National Academy of Sciences*, 121(4):e2320424121, 2024.
[29] L. Anselin. Local indicators of spatial association—lisa. *Geographical analysis*, 27(2):93–115, 1995.
[30] ANTARA. Indonesia safe from southeast asia heatwave: Brin researcher, 2024.
[31] E. Arcaute, E. Hatna, P. Ferguson, H. Youn, A. Johansson, and M. Batty. Constructing cities, deconstructing scaling laws. *Journal of the royal society interface*, 12(102):20140745, 2015.
[32] M. Arzaghi and J. V. Henderson. Networking off madison avenue. *The Review of Economic Studies*, 75(4):1011–1038, 2008.
[33] Asian Development Bank and BPS-Statistics Indonesia. The informal sector and informal employment in indonesia. Technical report, 2011.
[34] Asian Transport Observatory (ATO). Urban profiles. Technical report, United Nations Economic and Social Commission for Asia and the Pacific (UNESCAP), 2025.
[35] U. Aslak and L. Alessandretti. Infostop: scalable stop-location detection in multi-user mobility data. *arXiv preprint arXiv:2003.14370*, 2020.
[36] Associated Press. California work safety board approves indoor heat rules, but another state agency raises objections, 2024.
[37] Associated Press. Eastern half of u.s. sweltering again, with dangerous heat warnings across midwest and east, 2025.
[38] S. Athey, B. Ferguson, M. Gentzkow, and T. Schmidt. Estimating experienced racial segregation in us cities using large-scale gps data. *Proceedings of the National Academy of Sciences*, 118(46):e2026160118, 2021.
[39] S. Athey, B. Ferguson, M. Gentzkow, and T. Schmidt. Estimating experienced racial segregation in US cities using large-scale GPS data. *Proceedings of the National Academy of Sciences*, 118(46):e2026160118, 2021.
[40] S. Athey, B. Ferguson, M. Gentzkow, and T. Schmidt. Estimating experienced racial segregation in US cities using large-scale GPS data. *Proceedings of the National Academy of Sciences*, 118(46):e2026160118, Nov. 2021. Publisher: Proceedings of the National Academy of Sciences.
[41] D. Atkin, M. K. Chen, and A. Popov. The returns to face-to-face interactions: Knowledge spillovers in silicon valley. Technical report, National Bureau of Economic Research, 2022.
[42] R. Atkinson. Limited exposure: Social concealment, mobility and engagement with public space by the super-rich in london. *Environment and Planning A: Economy and Space*, 48(7):1302–1317, 2016.
[43] A. N. (Australia). New delhi's poor desperate for water as heatwave sparks 'tanker mafia', 2024.

[44] J. Bai, E. Brynjolfsson, W. Jin, S. Steffen, and C. Wan. Your company will need remote work as extreme weather gets worse. *Harvard Business Review*.
[45] M. Bailey, R. Cao, T. Kuchler, J. Stroebel, and A. Wong. Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, 32(3):259–280, 2018.
[46] M. Bailey, P. Farrell, T. Kuchler, and J. Stroebel. Social connectedness in urban areas. *Journal of Urban Economics*, 118:103264, 2020.
[47] J. Balaguer, H. Spiers, D. Hassabis, and C. Summerfield. Neural mechanisms of hierarchical planning in a virtual subway network. *Neuron*, 90(4):893–903, 2016.
[48] J. Ballester, M. Quijal-Zamorano, R. F. Méndez Turrubiates, F. Pegenaute, F. R. Herrmann, J. M. Robine, X. Basagaña, C. Tonne, J. M. Antó, and H. Achebak. Heat-related mortality in europe during the summer of 2022. *Nature medicine*, 29(7):1857–1866, 2023.
[49] H. Barbosa, M. Barthelemy, G. Ghoshal, C. R. James, M. Lenormand, T. Louail, R. Menezes, J. J. Ramasco, F. Simini, and M. Tomasini. Human mobility: Models and applications. *Physics Reports*, 734:1–74, 2018.
[50] A. Barreca, K. Clay, O. Deschenes, M. Greenstone, and J. S. Shapiro. 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, 2016.
[51] A. Barreca, K. Clay, O. Deschênes, M. Greenstone, and J. S. Shapiro. Adapting to climate change: The remarkable decline in the u.s. temperature–mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1):105–159, 2016.
[52] J. M. Barrero, N. Bloom, and S. J. Davis. Why working from home will stick. Technical report, National Bureau of Economic Research, 2021.
[53] J. M. Barrero, N. Bloom, and S. J. Davis. The evolution of work from home. Technical report, National Bureau of Economic Research, 2023.
[54] J. M. Barrero, N. Bloom, and S. J. Davis. Long social distancing. *Journal of Labor Economics*, 41(S1):S129–S172, 2023.
[55] R. Basu, N. Colaninno, A. Alhassan, and A. Sevtsuk. Hot and bothered: Exploring the effect of heat on pedestrian route choice behavior and accessibility. *Cities*, 155:105435, 2024.
[56] M. Batty. *Cities and Complexity: Understanding Cities with Cellular Automata, Agent-Based Models, and Fractals*. MIT Press, 2007.
[57] M. Batty. *The New Science of Cities*. MIT Press, 2013.
[58] I. Batur, V. O. Alhassan, M. V. Chester, S. E. Polzin, C. Chen, C. R. Bhat, and R. M. Pendyala. Understanding how extreme heat impacts human activity-mobility and time use patterns. *Transportation Research Part D: Transport and Environment*, 136:104431, 2024.
[59] N. Baum-Snow. Did highways cause suburbanization? *The quarterly journal of economics*, 122(2):775–805, 2007.
[60] BBC News. Luton airport flights suspended after runway surface defect found amid heatwave, 2022.
[61] R. Bean and et al. How does weather affect bikeshare use? a comparative analysis of forty cities across climate zones. *Journal of Transport Geography*, 95:103154, 2021.
[62] R. A. Beauregard and A. Haila. The unavoidable incompleteness of the city. *American Behavioral Scientist*, 41(3):327–341, 1997.
[63] H. E. Beck, N. E. Zimmermann, T. R. McVicar, N. Vergopolan, A. Berg, and E. F. Wood. Present and future köppen-geiger climate classification maps at 1-km resolution. *Scientific data*, 5(1):1–12, 2018.
[64] M. Beestermöller, L. Jessen-Thiesen, and A. Sandkamp. Striking evidence: the impact of railway strikes on competition from intercity bus services in germany. *Journal of Economic Geography*, page lbaf017, 2025.
[65] E. Berkes and R. Gaetani. The geography of unconventional innovation. *The Economic Journal*, 131(636):1466–1514, 2021.
[66] A. Bertaud. *Order without design: How markets shape cities*. The MIT Press, 2018.
[67] A. Bertaud. Order without design: How markets shape cities. *Town and Regional Planning*, 79:2–5, 2021.
[68] L. M. Bettencourt. The origins of scaling in cities. *science*, 340(6139):1438–1441, 2013.
[69] L. M. Bettencourt. Introduction to urban science: evidence and theory of cities as complex

systems. 2021.
[70] L. M. Bettencourt, J. Lobo, D. Helbing, C. Kühnert, and G. B. West. Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the national academy of sciences*, 104(17):7301–7306, 2007.
[71] L. M. A. Bettencourt, J. Lobo, D. Helbing, C. Kühnert, and G. B. West. Growth, innovation, scaling and the pace of life in cities. *PNAS*, 104(17):7301–7306, 2007.
[72] J. L. Blois, J. W. Williams, M. C. Fitzpatrick, S. T. Jackson, and S. Ferrier. Space can substitute for time in predicting climate-change effects on biodiversity. *Proceedings of the national academy of sciences*, 110(23):9374–9379, 2013.
[73] N. Bloom and A. Finan. The rise in super commuters. *Stanford Institute for Economic Research*, 2024.
[74] N. Bloom, C. I. Jones, J. Van Reenen, and M. Webb. Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144, 2020.
[75] N. Bloom, J. Liang, J. Roberts, and Z. J. Ying. Does working from home work? evidence from a chinese experiment. *Quarterly Journal of Economics*, 130(1):165–218, 2015.
[76] N. Bloom, M. Schankerman, and J. Van Reenen. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393, 2013.
[77] J. Blumgart. How redlining segregated philadelphia. *WHYY News*.
[78] N. Blüthgen, M. Staab, R. Achury, and W. W. Weisser. Unravelling insect declines: can space replace time? *Biology Letters*, 18(4):20210666, 2022.
[79] L. Böcker, M. Dijst, and J. Prillwitz. Impact of everyday weather on individual daily travel behaviours in perspective: a literature review. *Transport reviews*, 33(1):71–91, 2013.
[80] L. Böcker, T. P. Uteng, C. Liu, and M. Dijst. Weather and daily mobility in international perspective: A cross-comparison of dutch, norwegian and swedish city regions. *Transportation research part D: transport and environment*, 77:491–505, 2019.
[81] E. Bokányi, S. Juhász, M. Karsai, and B. Lengyel. Universal patterns of long-distance commuting and social assortativity in cities. *Scientific reports*, 11(1):20829, 2021.
[82] G. Bonaccorsi, F. Pierri, M. Cinelli, A. Flori, A. Galeazzi, F. Porcelli, A. L. Schmidt, C. M. Valensise, A. Scala, W. Quattrociocchi, et al. Economic and social consequences of human mobility restrictions under covid-19. *Proceedings of the National Academy of Sciences*, 117(27):15530–15535, 2020.
[83] C. Bongiorno, Y. Zhou, M. Kryven, D. Theurel, A. Rizzo, P. Santi, J. Tenenbaum, and C. Ratti. Vector-based pedestrian navigation in cities. *Nature Computational Science*, 1(10):678–685, 2021.
[84] V. Borràs, M. Ajenjo, and S. Moreno-Colom. More time parenting in spain: a possible change towards gender equality? *Journal of Family Studies*, 2021.
[85] L. Boucherie, B. F. Maier, and S. Lehmann. Decoupling geographical constraints from human mobility. *Nature Human Behaviour*, 2025. Online ahead of print.
[86] L. Breiman. *Classification and regression trees*. Routledge, 2017.
[87] C. R. Browning, C. A. Calder, L. J. Krivo, A. L. Smith, and B. Boettner. Socioeconomic segregation of activity spaces in urban neighborhoods: Does shared residence mean shared routines? *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 3(2):210–231, 2017.
[88] C. R. Browning, J. Tarrence, C. A. Calder, N. P. Pinchak, and B. Boettner. Geographic isolation, compelled mobility, and everyday exposure to neighborhood racial composition among urban youth. *American Journal of Sociology*, 128(3):914–961, 2022.
[89] K. Büchel and M. V. Ehrlich. Cities and the structure of social interactions: Evidence from mobile phone data. *Journal of urban economics*, 119:103276, 2020.
[90] K. Burkart, F. Meier, D. Scherer, A. Heidemann, W. R. Endlicher, M. Blüher, and J. Schmitt. Modification of heat-related mortality in an elderly urban population by vegetation (urban green) and proximity to water (urban blue): Evidence from lisbon, portugal. *Environmental Health Perspectives*, 124(7):927–934, 2016.
[91] M. Burke, S. M. Hsiang, and E. Miguel. Climate and conflict. *Annu. Rev. Econ.*, 7(1):577–617, 2015.
[92] M. Burke, S. M. Hsiang, and E. Miguel. Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239, 2015.
[93] M. Burke, S. M. Hsiang, and E. Miguel. Global non-linear effect of temperature on eco-

nomic production. *Nature*, 527:235–239, 2015.

[94] M. Burke, M. Zahid, M. C. Martins, C. W. Callahan, R. Lee, T. Avirmed, S. Heft-Neal, M. Kiang, S. M. Hsiang, and D. Lobell. Are we adapting to climate change? Technical report, National Bureau of Economic Research, 2024.

[95] C. Cabrera-Arnau, C. Zhong, M. Batty, R. Silva, and S. M. Kang. Inferring urban polycentricity from the variability in human mobility patterns. *Scientific Reports*, 13(1):5751, 2023.

[96] K. A. Cagney, E. York Cornwell, A. W. Goldman, and L. Cai. Urban mobility and activity space. *Annual Review of Sociology*, 46(1):623–648, 2020.

[97] K. A. Cagney, E. York Cornwell, A. W. Goldman, and L. Cai. Urban mobility and activity space. *Annual Review of Sociology*, 46:623–648, 2020.

[98] F. Calabrese and C. Ratti. An urban-wide real-time monitoring system: The real time rome project. *NETCOM: Réseaux, communication et territoires/Networks and communication studies*, 20(3):247–257, 2006.

[99] F. Calabrese, Z. Smoreda, V. D. Blondel, and C. Ratti. Interplay between telecommunications and face-to-face interactions: A study using mobile phone data. *PloS one*, 6(7):e20814, 2011.

[100] Cal/OSHA Standards Board. Indoor heat illness prevention standard, 2024. Board approval materials and final text.

[101] J. Candia, M. C. González, P. Wang, T. Schoenharl, G. Madey, and A.-L. Barabási. Uncovering individual and collective human dynamics from mobile phone records. *Journal of physics A: mathematical and theoretical*, 41(22):224015, 2008.

[102] J. Candipan, N. E. Phillips, R. J. Sampson, and M. Small. From residence to movement: The nature of racial segregation in everyday urban mobility. *Urban Studies*, 58(15):3095–3117, 2021.

[103] M. Castells. *The informational city: Information technology, economic restructuring, and the urban-regional process*, volume 1. Blackwell Oxford, 1989.

[104] CCOO Andalucía. Prevención ante olas de calor: paradas, jornadas intensivas y derechos laborales, 2025.

[105] Centre for Cities. Cities outlook 2024. `https://www.centreforcities.org/publication/cities-outlook-2024/`, 2024.

[106] S. Chang, E. Pierson, P. W. Koh, J. Gerardin, B. Redbird, D. Grusky, and J. Leskovec. Mobility network models of covid-19 explain inequities and inform reopening. *Nature*, 589(7840):82–87, 2021.

[107] K. Chatterjee, S. Chng, B. Clark, A. Davis, J. De Vos, D. Ettema, S. Handy, A. Martin, and L. Reardon. Commuting and wellbeing: a critical overview of the literature with implications for policy and future research. *Transport reviews*, 40(1):5–34, 2020.

[108] M. A. Chen and S. Sinha. Home-based workers and cities. *Environment and Urbanization*, 28(2):343–358, 2016.

[109] M. K. Chen and R. Rohla. The effect of partisanship and political advertising on close family ties. *Science*, 360(6392):1020–1024, 2018.

[110] Y. Chen, J. Ying, Y. Zhang, Y. Feng, X. Chen, and X. Zhou. A systematic review of urban green spaces affecting subjective well-being: an explanation of their complex mechanisms. *Ecological Indicators*, 178:113786, 2025.

[111] Z. Chen, S. Kelty, A. G. Evsukoff, B. F. Welles, J. Bagrow, R. Menezes, and G. Ghoshal. Contrasting social and non-social sources of predictability in human mobility. *Nature Communications*, 13(1):1922, Apr. 2022. Number: 1 Publisher: Nature Publishing Group.

[112] Z. Chen, B. Yu, W. Song, H. Liu, Q. Wu, K. Shi, and J. Wu. A new approach for detecting urban centers and their spatial structure with nighttime light remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 55(11):6305–6319, 2017.

[113] R. Chetty, W. S. Dobbie, B. Goldman, S. Porter, and C. Yang. Changing opportunity: Sociological mechanisms underlying growing class gaps and shrinking race gaps in economic mobility. Technical report, National Bureau of Economic Research, 2024.

[114] R. Chetty, J. N. Friedman, N. Hendren, M. Stepner, and the Opportunity Insights Team. The economic impacts of covid-19: Evidence from a new public database built using private sector data. Technical Report w27431, NBER, 2020.

[115] R. Chetty and N. Hendren. The impacts of neighborhoods on intergenerational mobility:

Childhood exposure effects. *The Quarterly Journal of Economics*, 133(3):1107–1162, 2018.

[116] R. Chetty, N. Hendren, and L. F. Katz. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902, 2016.

[117] R. Chetty, N. Hendren, P. Kline, and E. Saez. Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The quarterly journal of economics*, 129(4):1553–1623, 2014.

[118] R. Chetty, M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, et al. Social capital i: measurement and associations with economic mobility. *Nature*, 608(7921):108–121, 2022.

[119] R. Chetty, M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, et al. Social capital ii: determinants of economic connectedness. *Nature*, 608(7921):122–134, 2022.

[120] M. Chinazzi, J. T. Davis, M. Ajelli, C. Gioannini, M. Litvinova, S. Merler, A. Pastore y Piontti, K. Mu, L. Rossi, K. Sun, et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak. *Science*, 368(6489):395–400, 2020.

[121] J. Choi, J. Guzman, and M. L. Small. Third places and neighborhood entrepreneurship: Evidence from starbucks cafés. Technical report, National Bureau of Economic Research, 2024.

[122] W. Christaller and C. W. Baskin. Central places in southern germany. *(No Title)*, 1966.

[123] E. Chyn and L. F. Katz. Neighborhoods matter: Assessing the evidence for place effects. *Journal of Economic Perspectives*, 35(4):197–222, 2021.

[124] City of Phoenix. Heat relief network and cooling centers: Extended and 24/7 sites, 2025. Program overview and seasonal sites; includes extended-hour and overnight hub details.

[125] D. T. Coe and E. Helpman. International r&d spillovers. *European Economic Review*, 39(5):859–887, 1995.

[126] E. D. Coffel, T. R. Thompson, and R. M. Horton. The impacts of rising temperatures on aircraft takeoff performance. *Climatic Change*, 144(2):381–388, 2017.

[127] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *Proceedings of the National Academy of Sciences*, 102(21):7426–7431, 2005.

[128] R. Colacito, B. Hoffmann, and T. Phan. Temperature and growth: A panel analysis of the united states. *Journal of Money, Credit and Banking*, 51(2-3):313–368, 2019.

[129] P.-P. Combes, G. Duranton, and L. Gobillon. The productivity advantages of large cities: Distinguishing agglomeration from firm selection. *Econometrica*, 80(6):2543–2594, 2012.

[130] P.-P. Combes, G. Duranton, and L. Gobillon. The costs of agglomeration: House and land prices in french cities. *The Review of Economic Studies*, 86(4):1556–1589, 2019.

[131] S. Cong, D. Nock, Y. L. Qiu, and B. Xing. Unveiling hidden energy poverty using the energy equity gap. *Nature communications*, 13(1):2456, 2022.

[132] C. Cook, L. Currier, and E. Glaeser. Urban mobility and the experienced isolation of students. *Nature Cities*, 1(1):73–82, 2024.

[133] Copernicus Climate Change Service (C3S). ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate, 2017. Copernicus Climate Change Service Climate Data Store (CDS), accessed [Insert Date of Access].

[134] C. Cornish. Learning to live with 50c temperatures. *Financial Times*, August 2024. Accessed September 22, 2024.

[135] C. Cottineau. Metazipf. a dynamic meta-analysis of city size distributions. *PloS one*, 12(8):e0183919, 2017.

[136] C. Cottineau, O. Finance, E. Hatna, E. Arcaute, and M. Batty. Defining urban clusters to detect agglomeration economies. *Environment and Planning B: Urban Analytics and City Science*, 46(9):1611–1626, 2019.

[137] V. Couture, G. Duranton, and M. A. Turner. Speed. *Review of Economics and Statistics*, 100(4):725–739, 2018.

[138] V. Couture and J. Handbury. Urban revival in america. *Journal of Urban Economics*, 119:103267, 2020.

[139] P. F. Craigmile and P. Guttorp. Comparing cmip6 climate model simulations of annual

global mean temperatures to a new combined data product. *Earth and Space Science*, 10(10):e2022EA002468, 2023.

[140] K. Credit, O. Kekezi, C. Mellander, and R. Florida. Third places, the connective fibre of cities and high-tech entrepreneurship. *Regional studies*, 58(12):2225–2240, 2024.

[141] I. Crunchbase. Crunchbase api. `https://data.crunchbase.com/docs/using-the-api`, 2025.

[142] J.-L. Cruz. Global warming and labor market reallocation. *Working Paper*, 2021.

[143] Z. Cullen, B. Pakzad-Hurson, and R. Perez-Truglia. Home sweet home: How much do employees value remote work? In *AEA Papers and Proceedings*, volume 115, pages 276–281. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2025.

[144] D. M. Cutler and E. L. Glaeser. Are ghettos good or bad? *The Quarterly Journal of Economics*, 112(3):827–872, 1997.

[145] D. M. Cutler, E. L. Glaeser, and J. L. Vigdor. When are ghettos bad? lessons from immigrant segregation in the united states. *Journal of Urban Economics*, 63(3):759–774, 2008.

[146] R. M. Danner, C. M. Coomes, and E. P. Derryberry. Simulated heat waves reduce cognitive and motor performance of an endotherm. *Ecology and Evolution*, 11:2261–2272, 2021.

[147] C. Dauth, S. Findeisen, and J. Südekum. Sorting, selection and agglomeration effects across german cities. *Journal of Labor Economics*, 40(S1):S1–S50, 2022.

[148] G. Davies, J. Dixon, C. G. Tredoux, J. D. Whyatt, J. J. Huck, B. Sturgeon, B. T. Hocking, N. Jarman, and D. Bryan. Networks of (dis) connection: mobility practices, tertiary streets, and sectarian divisions in north belfast. *Annals of the American Association of Geographers*, 109(6):1729–1747, 2019.

[149] D. R. Davis, J. I. Dingel, J. Monras, and E. Morales. How Segregated Is Urban Consumption? *Journal of Political Economy*, 127(4):1684–1738, Aug. 2019. Publisher: The University of Chicago Press.

[150] D. R. Davis, J. I. Dingel, J. Monras, and E. Morales. How segregated is urban consumption? *Journal of Political Economy*, 127(4):1684–1738, 2019.

[151] L. W. Davis and P. J. Gertler. Contribution of air conditioning adoption to future energy use under global warming. *Proceedings of the National Academy of Sciences*, 112(19):5962–5967, 2015.

[152] S. J. Davis. The big shift in working arrangements: Eight ways unusual. Technical report, National Bureau of Economic Research, 2024.

[153] À. G. de la Prada and M. L. Small. How people are exposed to neighborhoods racially different from their own. *Proceedings of the National Academy of Sciences*, 121(28):e2401661121, 2024.

[154] G. De Rassenfosse, J. Kozak, and F. Seliger. Geocoding of worldwide patent data. *Scientific data*, 6(1):260, 2019.

[155] O. . E. S. de Sinaloa. Ambulantes reportan caída de 60% por el calor, 2024.

[156] E. Dean. A brief history of white people in southeast. *Washington City Paper*.

[157] M. Dell, B. F. Jones, and B. A. Olken. Temperature shocks and economic growth: Evidence from the last half century. *AEJ: Macroeconomics*, 4(3):66–95, 2012.

[158] M. Dell, B. F. Jones, and B. A. Olken. What do we learn from the weather? the new climate-economy literature. *Journal of Economic literature*, 52(3):740–798, 2014.

[159] M. J. Delventhal, S. Kwon, and A. Parkhomenko. How do cities change when we work from home? *Journal of Urban Economics*, 127:103437, 2022.

[160] Detik. Bmkg: Potensi suhu panas siang hari saat kemarau; imbauan pengurangan aktivitas luar ruang, 2024.

[161] C. Di Napoli, C. Barnard, C. Prudhomme, H. L. Cloke, and F. Pappenberger. Era5-heat: a global gridded historical dataset of human thermal comfort indices from climate reanalysis. *Geoscience data journal*, 8(1):2–10, 2021.

[162] J. I. Dingel and B. Neiman. How many jobs can be done at home? *Journal of Public Economics*, 189:104235, 2020.

[163] J. Doan. Validating spend data for brands against company reporting. `https://www.safegraph.com/blog/validating-spend-data-for-brands-against-company-reporting`, 12 2021.

Accessed: September 25, 2024.
[164] K. Dobney, C. J. Baker, A. Quinn, and L. Chapman. Quantifying the effects of high summer temperatures due to climate change on buckling and rail related delays in south-east united kingdom. *Meteorological Applications*, 16(2):245–251, 2009.
[165] L. Dong, F. Duarte, G. Duranton, P. Santi, M. Barthelemy, M. Batty, L. Bettencourt, M. Goodchild, G. Hack, Y. Liu, et al. Defining a city—delineating urban areas using cell-phone data. *Nature Cities*, 1(2):117–125, 2024.
[166] X. Dong, A. J. Morales, E. Jahani, E. Moro, B. Lepri, B. Bozkaya, C. Sarraute, Y. Bar-Yam, and A. Pentland. Segregated interactions in urban and online space. *EPJ Data Science*, 9(1):1–22, Dec. 2020. Number: 1 Publisher: SpringerOpen.
[167] R. I. Dunbar. Neocortex size as a constraint on group size in primates. *Journal of human evolution*, 22(6):469–493, 1992.
[168] M. Duneier. *Ghetto: The invention of a place, the history of an idea*. Macmillan, 2016.
[169] D. M. Dunlavy, T. G. Kolda, and E. Acar. Temporal link prediction using matrix and tensor factorizations. *ACM Transactions on Knowledge Discovery from Data*, 5(2):1–27, 2011.
[170] G. Duranton. Urban evolutions: The fast, the slow, and the still. *American Economic Review*, 97(1):197–221, 2007.
[171] G. Duranton and D. Puga. Micro-foundations of urban agglomeration economies. In *Handbook of regional and urban economics*, volume 4, pages 2063–2117. Elsevier, 2004.
[172] Y. Dzyuban, A. Middel, A. Jankovic, et al. Public transit infrastructure and heat perceptions in hot and dry climates. *International Journal of Biometeorology*, 66:179–192, 2022.
[173] F. Echenique and J. Fryer, Roland G. A measure of segregation based on social interactions. *The Quarterly Journal of Economics*, 122(2):441–485, 2007.
[174] El País.
[175] El País. Madrid activa una xarxa de refugis climàtics i espais amb aire condicionat durant les onades de calor, 2024.
[176] El País. Barcelona amplía la red de refugios climáticos a alrededor de 400 espacios, 2025.
[177] El País. Derechos laborales ante las olas de calor: cuándo se puede interrumpir la actividad por riesgo, 2025. Explainer on worker rights and heat alerts.
[178] El País. Madrid cierra parques y limita el acceso en episodios de calor extremo, 2025.
[179] T. (English). Bmkg cites causes of jakarta's recent hot days; advice on timing activity, 2024.
[180] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, pages 226–231, 1996.
[181] F. Estrada, W. W. Botzen, and R. S. Tol. A global economic assessment of city policies to reduce climate change impacts. *Nature Climate Change*, 7(6):403–406, 2017.
[182] EURopean Employment Services. Labour market information: Spain, 2024. Accessed: December 12, 2024.
[183] M. E. Evans, P. B. Adler, A. L. Angert, S. M. Dey, M. P. Girardin, K. A. Heilman, S. Klesse, D. L. Perret, D. F. Sax, S. N. Sheth, et al. Reconsidering space-for-time substitution in climate change ecology. *Nature Climate Change*, pages 1–4, 2025.
[184] Excélsior. La ola de calor altera hábitos de consumo (menos compras en calle; más en línea/ac), 2024.
[185] Z. Fan, T. Su, M. Sun, A. Noyman, F. Zhang, A. entland, and E. Moro. Diversity beyond density: Experienced social mixing of urban streets. *PNAS nexus*, 2(4):pgad077, 2023.
[186] E. Ferranti, L. Chapman, C. Lowe, S. McCulloch, D. Jaroszweski, and A. Quinn. Heat-related failures on southeast england's railway network: Insights and implications for heat risk management. *Weather, Climate, and Society*, 8(2):177–191, 2016.
[187] K. Finkel. Roots of hypersegregation in philadelphia, 1920–1930. *PhillyHistory Blog*.
[188] R. Florida. The world is spiky. *The Atlantic Monthly*, pages 48–51, Oct. 2005.
[189] R. Florida. *The new urban crisis: Gentrification, housing bubbles, growing inequality, and what we can do about it*. Simon and Schuster, 2017.
[190] R. Florida. *The rise of the creative class*. Basic books, 2019.
[191] C. for International Earth Science Information Network CIESIN Columbia University. Global gridded relative deprivation index (grdi), version 1, 2022.
[192] G. B. Ford. The city scientific. *Engineering Record*, 67:551–552, 1913.
[193] I. Foursquare Labs. Foursquare open source places (fsq os places): Global points of in-

terest dataset. https://opensource.foursquare.com/os-places/, 2024. Contains 100M+ global POIs under Apache License 2.0; accessed: YYYY-MM-DD.
[194] M. Fujita, P. R. Krugman, and A. Venables. *The spatial economy: Cities, regions, and international trade*. MIT press, 2001.
[195] C. Gaigné, H. R. Koster, F. Moizeau, and J.-F. Thisse. Who lives where in the city? amenities, commuting and income sorting. *Journal of Urban Economics*, 128:103394, 2022.
[196] E. Gallo, M. Quijal-Zamorano, R. F. Méndez Turrubiates, C. Tonne, X. Basagaña, H. Achebak, and J. Ballester. Heat-related mortality in europe during 2023 and the role of adaptation in protecting health. *Nature medicine*, pages 1–5, 2024.
[197] R. Gallotti, D. Maniscalco, M. Barthelemy, and M. De Domenico. The distorting lens of human mobility data. *arXiv preprint arXiv:2211.10308*, 2022.
[198] A. Getis and J. Getis. Christaller's central place theory. *Journal of Geography*, 65(5):220–226, 1966.
[199] J. J. Gibson. The theory of affordances:(1979). In *The people, place, and space reader*, pages 56–60. Routledge, 2014.
[200] W. Gibson. The future is already here—it's just not evenly distributed, 1992. Quotation, attributed to interviews and public talks by William Gibson.
[201] G. Gigerenzer and W. Gaissmaier. Heuristic decision making. *Annual review of psychology*, 62(2011):451–482, 2011.
[202] E. Glaeser. *Triumph of the city: How urban spaces make us human*. Pan Macmillan, 2011.
[203] E. L. Glaeser, J. Kolko, and A. Saiz. Consumer city. *Journal of economic geography*, 1(1):27–50, 2001.
[204] A. Goldstein, A. Kapelner, J. Bleich, and E. Pitkin. Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *journal of Computational and Graphical Statistics*, 24(1):44–65, 2015.
[205] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *nature*, 453(7196):779–782, 2008.
[206] M. González-Leonardo, C. Cabrera-Arnau, and F. Rowe. Harnessing digital data to strengthen disaster response in valencia's 2024 floods. *Population Europe*, 2024.
[207] Google. Covid-19 community mobility reports, 2022. Final update Oct 2022; Accessed 2025-09-23.
[208] C. Gorback. Ridesharing and the redistribution of economic activity. *Available at SSRN 4956559*, 2022.
[209] N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 2017.
[210] J. Graff Zivin and M. Neidell. Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26, 2014.
[211] M. S. Granovetter. The strength of weak ties. *American journal of sociology*, 78(6):1360–1380, 1973.
[212] B. N. Graphics. What the pret index told us about the economic recovery, 2022. Accessed 2025-09-23.
[213] T. L. Griffin. South side land narratives: The lost histories and hidden joys of black chicago. *Harvard Design Magazine*, (49), 2021.
[214] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864, 2016.
[215] Grupo 2000. Jornada intensiva de verano 2025: regulación laboral y acuerdos sectoriales, 2025.
[216] X. Gu, P. Chen, and C. Fan. Socio-demographic inequalities in the impacts of extreme temperatures on population mobility. *Journal of Transport Geography*, 114:103755, 2024.
[217] A. Gupta, V. Mittal, and S. Van Nieuwerburgh. Work from home and the office real estate apocalypse. Technical report, National Bureau of Economic Research, 2022.
[218] J. Gyourko, C. Mayer, and T. Sinai. Superstar cities. *American Economic Journal: Economic Policy*, 5(4):167–199, 2013.
[219] M. Hahsler, M. Piekenbrock, and D. Doran. dbscan: Fast density-based clustering with r. *Journal of Statistical Software*, 91:1–30, 2019.

[220] S. Hajat, M. O'Connor, and T. Kosatsky. Health effects of hot weather: from awareness of risk factors to effective health protection. *The Lancet*, 375(9717):856–863, 2010.
[221] S. Hallegatte, C. Green, R. J. Nicholls, and J. Corfee-Morlot. Future flood losses in major coastal cities. *Nature Climate Change*, 3(9):802–806, 2013.
[222] M. Harari. Cities in bad shape: Urban geometry in india. *American Economic Review*, 110(8):2377–2421, 2020.
[223] L. P. Hartley. *The Go-Between*. Hamish Hamilton, London, 1953.
[224] B. J. Hatchett, T. Benmarhnia, K. Guirguis, K. VanderMolen, A. Gershunov, H. Kerwin, A. Khlystov, K. M. Lambrecht, and V. Samburova. Mobility data to aid assessment of human responses to extreme environmental conditions. *The Lancet Planetary Health*, 5(10):e665–e667, 2021.
[225] A. K. Heaney et al. Estimated impacts of ambient temperatures on bikeshare trips in new york city. *Environmental Health Perspectives*, 127(3):037001, 2019.
[226] S. Heblich, S. J. Redding, and D. M. Sturm. The making of the modern metropolis: evidence from london. *The Quarterly Journal of Economics*, 135(4):2059–2133, 2020.
[227] D. Herald. 50% of delhi street vendors report significant income loss due to extreme heat, 2024.
[228] J. J. Hess, N. A. Errett, G. McGregor, T. Busch Isaksen, Z. S. Wettstein, S. K. Wheat, and K. L. Ebi. Public health preparedness for extreme heat events. *Annual Review of Public Health*, 44(1):301–321, 2023.
[229] G. Heutel, N. H. Miller, and D. Molitor. Adaptation and the mortality effects of temperature across us climate regions. *Review of Economics and Statistics*, 103(4):740–753, 2021.
[230] B. Hillier, A. Penn, J. Hanson, T. Grajewski, and J. Xu. Natural movement: or, configuration and attraction in urban pedestrian movement. *Environment and Planning B: Planning and Design*, 20(1):29–66, 1993.
[231] R. M. Hilman, G. Iñiguez, and M. Karsai. Socioeconomic biases in urban mixing patterns of us metropolitan areas. *EPJ data science*, 11(1):32, 2022.
[232] J. Ho. Is your spend data geographically representative? `https://colab.research.google.com/drive/16eBT1dqnUK76GnBMgPMY5rJO4vBD2gdQ`, 12 2021. Accessed: September 25, 2024.
[233] D. Hoornweg and K. Pope. Population predictions for the world's largest cities in the 21st century. *Environment and urbanization*, 29(1):195–216, 2017.
[234] T. Horanont, S. Phithakkitnukoon, T. W. Leong, Y. Sekimoto, and R. Shibasaki. Weather effects on the patterns of people's everyday activities: A study using gps traces of mobile phone users. *PLOS ONE*, 8(12):e81153, 2013.
[235] C.-T. Hsieh, E. Hurst, C. I. Jones, and P. J. Klenow. The allocation of talent and u.s. economic growth. *Econometrica*, 87(5):1439–1474, 2019.
[236] X. Huang, Y. Jiang, and A. Mostafavi. The emergence of urban heat traps and human mobility in 20 us cities. *npj Urban Sustainability*, 4:1–10, 2024.
[237] R. J. Hyndman and Y. Khandakar. Automatic time series forecasting: the forecast package for r. *Journal of statistical software*, 27:1–22, 2008.
[238] IESR. Indonesia faces hot temperatures: Bmkg says not a heatwave; exposure risks remain, 2024.
[239] International Energy Agency. The future of cooling. Technical report, International Energy Agency, 2018. Accessed on November 29, 2024.
[240] International Energy Agency. The future of cooling: Opportunities for energy-efficient air conditioning, 2018. Global AC adoption and efficiency scenarios.
[241] International Labour Organization. Women and men in the informal economy: A statistical picture (3rd ed.), 2018.
[242] J. Jacobs. *The economy of cities*. Vintage, 2016.
[243] E. Jahani, S. P. Fraiberger, M. Bailey, and D. Eckles. Long ties, disruptive life events, and economic prosperity. *Proceedings of the National Academy of Sciences*, 120(28):e2211062120, 2023.
[244] O. Jay, A. Capon, P. Berry, C. Broderick, R. de Dear, G. Havenith, Y. Honda, R. S. Kovats, W. Ma, A. Malik, et al. Reducing the health effects of hot weather and heat extremes: from personal cooling strategies to green cities. *The Lancet*, 398(10301):709–724, 2021.
[245] M. E. Kahn. *Adapting to climate change*. Yale University Press, 2021.

[246] D. Kahneman, J. L. Knetsch, and R. H. Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6):1325–1348, 1990.
[247] D. Kahneman, A. B. Krueger, D. A. Schkade, N. Schwarz, and A. A. Stone. A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306(5702):1776–1780, 2004.
[248] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
[249] R. Kaiser, A. Le Tertre, J. Schwartz, C. A. Gotway, W. R. Daley, and C. H. Rubin. The effect of the 1995 heat wave in chicago on all-cause and cause-specific mortality. *American Journal of Public Health*, 97:S158–S162, 2007.
[250] Kastle Systems. Getting america back to work: Workplace occupancy barometer, 2020. Keycard/fob/app entries; 10-city barometer; Accessed 2025-09-23.
[251] S. M. Kaufman, C. Qing, N. Levenson, M. Hanson, et al. Transportation during and after hurricane sandy. 2012.
[252] W. Keller. International technology diffusion. *Journal of economic literature*, 42(3):752–782, 2004.
[253] G. P. Kenny, J. Yardley, C. Brown, R. J. Sigal, and O. Jay. Heat stress in older individuals and patients with common chronic diseases. *Cmaj*, 182(10):1053–1060, 2010.
[254] A. Kerrion. Containers: 2024 ranking of the world's major ports, April 2025. Reports Top 20 ports handled 414.6 million TEUs in 2024; implies $\tilde{5}0\%$ of global container throughput.
[255] M. Keuschnigg, S. Mutgan, and P. Hedström. Urban scaling and the regional divide. *Science advances*, 5(1):eaav0042, 2019.
[256] M. A. Killingsworth and D. T. Gilbert. A wandering mind is an unhappy mind. *Science*, 330(6006):932–932, 2010.
[257] K. Kim. Investigation on the effects of weather and calendar events on bike-sharing according to station trip patterns. *Journal of Transport Geography*, 66:309–320, 2018.
[258] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.
[259] J. P. Kossin, T. L. Olander, and K. R. Knapp. Global increase in major tropical cyclone exceedance probability over the past four decades. *PNAS*, 117(22):11975–11980, 2020.
[260] E. Kotov, R. Lovelace, and E. Vidal-Tortosa. *spanishoddata*, 2024.
[261] Á. J. Kovács, S. Juhász, E. Bokányi, and B. Lengyel. Income-related spatial concentration of individual social capital in cities. *Environment and Planning B: Urban Analytics and City Science*, 50(4):1072–1086, 2023.
[262] G. Kreindler. Peak-hour road congestion pricing: Experimental evidence and equilibrium implications. *Econometrica*, 92(4):1233–1268, 2024.
[263] E. Kumakura, Y. Ashie, and T. Ueno. Assessing the impact of summer heat on the movement of people in tokyo based on mobile phone location data. *Building and Environment*, 265:111952, 2024.
[264] M. Kummu, M. Taka, and J. H. Guillaume. Gridded global datasets for gross domestic product and human development index over 1990–2015. *Scientific data*, 5(1):1–15, 2018.
[265] Lancet Countdown. Heat-related Mortality. `https://lancetcountdown.org/explore-our-data/`, 2023.
[266] S. Larcom, F. Rauch, and T. Willems. The benefits of forced experimentation: Striking evidence from the london underground network. *The Quarterly Journal of Economics*, 132(4):2019–2055, 2017.
[267] J. G. C. Laurent, A. Williams, Y. Oulhote, A. Zanobetti, J. G. Allen, and J. D. Spengler. Reduced cognitive function during a heat wave among residents of non-air-conditioned buildings: An observational study of young adults in the summer of 2016. *PLoS medicine*, 15(7):e1002605, 2018.
[268] G. Le Roux, J. Vallée, and H. Commenges. Social segregation around the clock in the paris region (france). *Journal of Transport Geography*, 59:134–145, 2017.
[269] W. Lei, L. Jiao, Z. Xu, Z. Zhou, and G. Xu. Scaling of urban economic outputs: Insights both from urban population size and population mobility. *Computers, Environment and Urban Systems*, 88:101657, 2021.
[270] M. Leonardi and E. Moretti. The agglomeration of urban amenities: Evidence from milan restaurants. *American Economic Review: Insights*, 5(2):141–157, 2023.

[271] B. L. Levy, N. E. Phillips, and R. J. Sampson. Triple disadvantage: neighborhood networks of everyday urban mobility and violence in us cities. *American Sociological Review*, 85(6):925–956, 2020.
[272] S. Leydon. How a long-ago map created racial boundaries that still define boston. *GBH News (Boston)*.
[273] C. Li, G. Chen, and S. Wang. Urban mobility resilience under heat extremes: Evidence from bike-sharing travel in new york. *Travel Behaviour and Society*, 37:100821, 2024.
[274] H. Li, H. Campbell, and S. Fernandez. Residential segregation, spatial mismatch and economic growth across us metropolitan areas. *Urban Studies*, 50(13):2642–2660, 2013.
[275] W. Li, Q. Wang, Y. Liu, M. L. Small, and J. Gao. A spatiotemporal decay model of human mobility when facing large-scale crises. *Proceedings of the National Academy of Sciences*, 119(24):e2203042119, 2022.
[276] Z. Li, H. Ning, F. Jing, and M. N. Lessani. Understanding the bias of mobile location data across spatial scales and over time: a comprehensive analysis of safegraph data in the united states. *Plos one*, 19(1):e0294430, 2024.
[277] Y. Liao, J. Gil, S. Yeh, R. H. Pereira, and L. Alessandretti. Socio-spatial segregation and human mobility: A review of empirical evidence. *arXiv preprint arXiv:2403.06641*, 2024.
[278] Y. Liao, J. Gil, S. Yeh, R. H. Pereira, and L. Alessandretti. Socio-spatial segregation and human mobility: A review of empirical evidence. *Computers, Environment and Urban Systems*, 117:102250, 2025.
[279] Lincoln Institute of Land Policy. Fiscally standardized cities database, 2022.
[280] J. d. J. López Salazar and J. Vanek. Informal workers in mexico: A statistical snapshot. Technical report, WIEGO Statistical Brief No. 22, 2020.
[281] N. Lorenzo, A. Díaz-Poso, and D. Royé. Heatwave intensity on the iberian peninsula: Future climate projections. *Atmospheric Research*, 258:105655, 2021.
[282] A. Losch et al. Economics of location. 1954.
[283] T. Louail, M. Lenormand, O. G. Cantu Ros, M. Picornell, R. Herranz, E. Frias-Martinez, J. J. Ramasco, and M. Barthelemy. From mobile phone data to the spatial structure of cities. *Scientific reports*, 4(1):5276, 2014.
[284] T. Louail, M. Lenormand, J. Murillo Arias, and J. J. Ramasco. Crowdsourcing the Robin Hood effect in cities. *Applied Network Science*, 2(1):1–13, Dec. 2017. Number: 1 Publisher: SpringerOpen.
[285] T. Louail, M. Lenormand, J. Murillo Arias, and J. J. Ramasco. Crowdsourcing the robin hood effect in cities. *Applied network science*, 2:1–13, 2017.
[286] T. Louail, M. Lenormand, M. Picornell, O. Garcia Cantu, R. Herranz, E. Frias-Martinez, J. J. Ramasco, and M. Barthelemy. Uncovering the spatial structure of mobility networks. *Nature communications*, 6(1):6007, 2015.
[287] R. Louf and M. Barthelemy. Modeling the polycentric transition of cities. *Physical review letters*, 111(19):198702, 2013.
[288] R. Louf and M. Barthelemy. How congestion shapes cities: from mobility patterns to scaling. *Scientific Reports*, 4(1):5561, July 2014. Number: 1 Publisher: Nature Publishing Group.
[289] R. E. Lucas Jr. Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. North-Holland, 1976.
[290] S. Lüthi, C. Fairless, E. M. Fischer, N. Scovronick, B. Armstrong, M. D. S. Z. S. Coelho, Y. L. Guo, Y. Guo, Y. Honda, V. Huber, et al. Rapid increase in the risk of heat-related mortality. *Nature communications*, 14(1):4894, 2023.
[291] A. Ly, F. V. Davenport, and N. S. Diffenbaugh. Exploring the influence of summer temperature on human mobility during the covid-19 pandemic in the san francisco bay area. *GeoHealth*, 7:e2022GH000772, 2023.
[292] A. Ly and et al. Exploring the influence of summer temperature on human mobility in us metro areas. *International Journal of Environmental Research and Public Health*, 20(9):5730, 2023.
[293] S. Lychagin. Spillovers, absorptive capacity and agglomeration. *Journal of Urban Economics*, 96:17–35, 2016.
[294] A. Mahajan. Highways and segregation. *Journal of Urban Economics*, 141:103574, 2024.
[295] C. Marchetti. Anthropological invariants in travel behavior. *Technological forecasting and*

social change, 47(1):75–88, 1994.

[296] P. V. Marsden. Core discussion networks of americans. *American sociological review*, pages 122–131, 1987.

[297] A. Marshall. Principles of economics. *London: MacMillan*, pages 1–627, 1890.

[298] P. Masselot, M. Mistry, J. Vanoli, R. Schneider, T. Iungman, D. Garcia-Leon, J.-C. Ciscar, L. Feyen, H. Orru, A. Urban, et al. Excess mortality attributed to heat and cold: a health impact assessment study in 854 cities in europe. *The Lancet Planetary Health*, 7(4):e271–e281, 2023.

[299] M. Massenkoff and N. Wilmers. Rubbing shoulders: Class segregation in daily activities. *Available at SSRN 4516850*, 2023.

[300] D. S. Massey. The age of extremes: Concentrated affluence and poverty in the twenty-first century. *Demography*, 33:395–412, 1996.

[301] D. S. Massey and N. A. Denton. Hypersegregation in us metropolitan areas: Black and hispanic segregation along five dimensions. *Demography*, 26:373–391, 1989.

[302] D. S. Massey and N. A. Denton. American apartheid: Segregation and the making of the underclass. In *Social Stratification, Class, Race, and Gender in Sociological Perspective, Second Edition*, pages 660–670. Routledge, 2019.

[303] C. Matias and V. Miele. Statistical clustering of temporal networks through a dynamic stochastic block model. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1119–1141, 2017.

[304] A. I. McLeod. Kendall rank correlation and mann-kendall trend test. *R package Kendall*, 602:1–10, 2005.

[305] A. J. McMichael and E. Lindgren. Climate change: present and future risks to health, and necessary responses. *Journal of internal medicine*, 270(5):401–413, 2011.

[306] D. P. McMillen and S. C. Smith. The number of subcenters in large urban areas. *Journal of urban economics*, 53(3):321–338, 2003.

[307] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.

[308] L. Medina and F. Schneider. Shadow economies around the world: What did we learn over the last 20 years? Technical report, World Bank Policy Research Working Paper 8550, 2018.

[309] M. MELCHIORRI, R. I. MARI, P. FLORIO, M. SCHIAVINA, K. KRASNODEBSKA, P. POLITIS, J. UHL, M. PESARESI, L. MAFFENINI, P. SULIS, et al. Stats in the city–the ghsl urban centre database 2025.

[310] A. Metzger, Y. Baharav, L. Nichols, M. Finke, B. Saunders, P. Mitchell, G. A. Wellenius, K. Baughman McLeod, and K. Shickman. Beliefs and behaviors associated with the first named heat wave in seville spain 2022. *Scientific reports*, 14(1):9055, 2024.

[311] Q. Miao, E. W. Welch, and P. S. Sriraj. Extreme weather, public transport ridership and moderating effect of bus stop shelters. *Journal of Transport Geography*, 74:125–133, 2019.

[312] S. Milgram. The experience of living in cities: A psychological analysis. In *Annual Meeting of the American Psychological Association., Sep, 1969, Washington, DC, US; This paper is based on an Invited Address presented to the Division of General Psychology at the aforementioned meeting.* American Psychological Association, 1970.

[313] Ministerio de Sanidad. Plan nacional de actuaciones preventivas por altas temperaturas 2025, 2025.

[314] Ministry of Transport, Mobility and Urban Agenda of Spain, MITMA. Open Data Movilidad 2020-2021. `https://www.transportes.gob.es/ministerio/proyectos-singulares/estudios-de-movilidad-con-big-data/opendata-movilidad`, 2024.

[315] P. S. Mishra, S. Chopra, and I. Stamatopoulos. Fast-food stores with a drive-through recovered post-pandemic; stores without did not. *Stores without did not (April 11, 2024)*, 2024.

[316] Y. Miyauchi, K. Nakajima, and S. J. Redding. The economics of spatial mobility: Theory and evidence using smartphone data. Technical report, National Bureau of Economic Research, 2021.

[317] C. Molinero and S. Thurner. How the geometry of cities determines urban scaling laws. *Journal of the Royal Society interface*, 18(176):20200705, 2021.

[318] A. Mooney, C. Hodgson, I. Smith, and A. Williams. How an era of extreme heat is reshaping economies. *Financial Times*, 2023.
[319] E. Moretti. *The New Geography of Jobs*. Houghton Mifflin Harcourt, 2012.
[320] E. Moretti. The effect of high-tech clusters on the productivity of top inventors. *American Economic Review*, 111(10):3328–3375, 2021.
[321] E. Moro, D. Calacci, X. Dong, and A. Pentland. Mobility patterns are associated with experienced income segregation in large us cities. *Nature communications*, 12(1):1–10, 2021.
[322] E. Moro, D. Calacci, X. Dong, and A. Pentland. Mobility patterns are associated with experienced income segregation in large US cities. *Nature Communications*, 12(1):4633, July 2021. Number: 1 Publisher: Nature Publishing Group.
[323] M. L. Moss and C. Qing. The emergence of the super-commuter. Technical report, Rudin Center for Transportation Policy & Management, NYU Wagner, 2012.
[324] B. Moya-Gómez, M. Stępniak, J. C. García-Palomares, E. Frías-Martínez, and J. Gutiérrez. Exploring night and day socio-spatial segregation based on mobile phone data: The case of medellin (colombia). *Computers, Environment and Urban Systems*, 89:101675, 2021.
[325] R. Muth. *Cities and Housing*. University of Chicago Press, 1969.
[326] National Rail Enquiries. Hot weather and the railway, 2025.
[327] T. Naushirvanov, E. Elejalde, K. Kalimeri, E. Omodei, M. Karsai, and L. Ferres. Evacuation patterns and socioeconomic stratification in the context of wildfires. *EPJ Data Science*, 14(1):23, 2025.
[328] A. Nazish, K. Abbas, and E. Sattar. Health impact of urban green spaces: a systematic review of heat-related morbidity and mortality. *BMJ open*, 14(9):e081632, 2024.
[329] A. A. Neath and J. E. Cavanaugh. The bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(2):199–203, 2012.
[330] V. M. Netto, M. P. Soares, and R. Paschoalino. Segregated networks in the city. *International Journal of Urban and Regional Research*, 39(6):1084–1102, 2015.
[331] Network Rail. How we prepare the railway for hot weather, 2025.
[332] M. E. J. Newman. Mixing patterns in networks. *Physical Review E*, 67(2):026126, 2003.
[333] K. News. Phoenix council approves purchase of former school for year-round heat relief site, 2025.
[334] N. S. Ngo. Urban bus ridership, income, and extreme weather events. *Transportation Research Part D: Transport and Environment*, 77:464–475, 2019.
[335] H. Nilforoshan, W. Looi, E. Pierson, B. Villanueva, N. Fishman, Y. Chen, J. Sholar, B. Redbird, D. Grusky, and J. Leskovec. Human mobility networks reveal increased segregation in large cities. *Nature*, 624(7992):586–592, 2023.
[336] A. Noulas, S. Scellato, R. Lambiotte, M. Pontil, and C. Mascolo. A tale of many cities: Universal patterns in human urban mobility. *PLOS ONE*, 7(5):e37027, 2012.
[337] M. Novik. Greggs shares tumble as hot weather prompts profit warning. *Financial Times*, 2025.
[338] Occupational Safety and Health Administration. Heat injury and illness prevention in outdoor and indoor work settings: Rulemaking, 2025.
[339] OECD. Time use for work, care and other day-to-day activities, 2016. Accessed on December 30, 2024.
[340] OECD. Usual working hours per week by gender, 2024. Accessed on December 30, 2024.
[341] T. of India. Delhi metro records 60.17 lakh passenger journeys in may heat, 2024.
[342] R. Oldenburg. "the problem of place in america": from the great good place (1989). In *The urban design reader*, pages 285–295. Routledge, 2013.
[343] R. Oldenburg and D. Brissett. The third place. *Qualitative sociology*, 5(4):265–284, 1982.
[344] OpenTable. State of the industry: Seated diners (yoy), n.d. Accessed 2025-09-23.
[345] W. H. Organization. The challenge of obesity. `https://www.who.int/europe/news-room/fact-sheets/item/the-challenge-of-obesity`, 2022. Accessed: [Current Date].
[346] W. H. Organization. Diabetes. `https://www.who.int/europe/news-room/fact-sheets/item/diabetes`, 2023. Accessed: 2023-12-19.
[347] D. Oudin Åström, F. Bertil, and J. Rocklöv. Heat wave impact on morbidity and mortality

in the elderly population: A review of recent studies. *Maturitas*, 69(2):99–105, 2011.
[348] A. Owens. Building inequality: Housing segregation and income segregation. *Sociological Science*, 6:497–525, 2019.
[349] B. Pandey, C. Brelsford, and K. C. Seto. Rising infrastructure inequalities accompany urbanization and economic development. *Nature Communications*, 16(1):1193, 2025.
[350] L. Pappalardo, E. Manley, V. Sekara, and L. Alessandretti. Future directions in human mobility science. *Nature computational science*, 3(7):588–600, 2023.
[351] R. J. Park, J. Goodman, M. Hurwitz, and J. Smith. Heat and learning. *American Economic Journal: Economic Policy*, 12(2):306–39, 2020.
[352] Y. M. Park and M.-P. Kwan. Beyond residential segregation: A spatiotemporal approach to examining multi-contextual segregation. *Computers, Environment and Urban Systems*, 71:98–108, 2018.
[353] F. Pavanello, R. Pacelli, N. Berta, M. Schiavina, and E. De Cian. Air-conditioning and the adaptation cooling deficit in the global south. *Nature Communications*, 12:6163, 2021.
[354] E. País. El récord de calor abrasa a los vendedores ambulantes de ciudad de méxico, 2024.
[355] T. P. Peixoto. Nonparametric Bayesian inference of the microcanonical stochastic block model. *Physical Review E*, 95(1):012317, 2017.
[356] J. D. Périard, D. DeGroot, and O. Jay. Exertional heat stroke in sport and the military: epidemiology and mitigation. *Experimental physiology*, 107(10):1111–1121, 2022.
[357] S. E. Perkins, L. Alexander, and J. Nairn. Increasing frequency, intensity and duration of observed global heatwaves and warm spells. *Geophysical Research Letters*, 39(20), 2012.
[358] S. E. Perkins and L. V. Alexander. On the measurement of heat waves. *Journal of climate*, 26(13):4500–4517, 2013.
[359] S. Perkins-Kirkpatrick and S. Lewis. Increasing trends in regional heatwaves. *Nature communications*, 11(1):3357, 2020.
[360] T. F. Pettigrew. Intergroup contact theory. *Annual Review of Psychology*, 49:65–85, 1998.
[361] T. F. Pettigrew and L. R. Tropp. A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5):751–783, 2006.
[362] K. W. Phillips, K. A. Liljenquist, and M. A. Neale. Is the pain worth the gain? the advantages and liabilities of agreeing with socially distinct newcomers. *Personality and Social Psychology Bulletin*, 35(3):336–350, 2009.
[363] E. C. Pielou. The measurement of diversity in different types of biological collections. *Journal of Theoretical Biology*, 13:131–144, 1966.
[364] A. Político. Por ola de calor, estados ajustan o suspenden jornada escolar, 2024.
[365] T. J. Post. Indonesia sees hottest april in 40 years, 2024.
[366] T. J. Post. Outdoor workers at risk amid scorching heat, 2024.
[367] A. Press. Heat forces greek authorities to shut down acropolis during afternoon hours, 2024.
[368] A. Press. Mexico's drought and heatwave lead to ice rationing; water shortages, 2024.
[369] A. Press. Phoenix keeps some cooling stations open overnight; expands heat response, 2024.
[370] A. Press. Across the us, cities combine shade and education to help people beat the heat, 2025.
[371] A. Press. Greece imposes work breaks as a heat wave grips the country, 2025.
[372] T. Print. Delhi's water 'tanker mafia' surges amid heat and shortages, 2024.
[373] R. Prud'homme and C.-W. Lee. Size, sprawl, speed and the efficiency of cities. *Urban Studies*, 36(11):1849–1858, 1999.
[374] R. D. Putnam and E. P. Unum. Diversity and community in the 21st century: The 2006 johan skytte prize lecture. *Scandinavian Political Studies*, 30(2):137–174, 2007.
[375] M. Pyron. Characterizing communities. *Nature Education Knowledge*, 3(10):39, 2010.
[376] Y. Qiao, J. Santos, A. M. K. Stoner, and G. Flintsch. Climate change impacts on asphalt road pavement construction and maintenance: an economic life-cycle assessment of adaptation measures in virginia (us). *Journal of Industrial Ecology*, 24(2):342–355, 2020.
[377] G. Quach. Delivery workers feel the heat of climate change. *Financial Times*, 2023.
[378] Quadratín. Vendedores ambulantes afectados por el calor (ventas hasta 60%), 2024.
[379] A. Ramani, J. Alcedo, and N. Bloom. How working from home reshapes cities. *Proceedings of the National Academy of Sciences*, 121(45):e2408930121, 2024.

[380] A. Ramani and N. Bloom. The donut effect of covid-19 on cities. Technical report, National Bureau of Economic Research, 2021.
[381] M. A. Rapino and A. E. Fields. Mega commuting in the u.s.: Time and distance in defining the long commute using the american community survey. Technical Report SEHSD-WP2013-03, U.S. Census Bureau, 2013.
[382] C. Ratti, S. Sobolevsky, F. Calabrese, C. Andris, J. Reades, M. Martino, R. Claxton, and S. H. Strogatz. Redrawing the map of great britain from a network of human interactions. *PloS one*, 5(12):e14248, 2010.
[383] G. Raveendran. Home-based workers in india: A statistical profile. Technical report, WIEGO Statistical Brief No. 23, 2020.
[384] S. F. Reardon and K. Bischoff. Income inequality and income segregation. *American Journal of Sociology*, 116(4):1092–1153, 2011.
[385] S. F. Reardon and G. Firebaugh. Measures of multigroup segregation. *Sociological Methodology*, 32:33–67, 2002.
[386] S. J. Redding and D. M. Sturm. The costs of remoteness: Evidence from german division and reunification. *American Economic Review*, 98(5):1766–1797, 2008.
[387] Y. Ren, M. Ercsey-Ravasz, P. Wang, M. C. González, and Z. Toroczkai. Predicting commuter flows in spatial networks using a radiation model based on temporal ranges. *Nature Communications*, 5:5347, 2014.
[388] Renfe. Plan de contingencia por altas temperaturas: medidas de servicio y suministro de agua, 2024.
[389] Reuters. Heat rule in play as temperature soars at u.s. open, 2018.
[390] Reuters. Why uk trains struggle in the heat: speed restrictions and buckling risk, 2022.
[391] Reuters. Italy issues sectoral measures and limits outdoor work amid severe heat, 2023.
[392] Reuters. Low rhine water hampers shipping; barges cut loads and add surcharges, 2023.
[393] Reuters. Qatar can lead gulf on protections for workers against heat stress, 2023.
[394] Reuters. Spain to ban some outdoor working during extreme heat conditions, 2023.
[395] Reuters. Ups agrees to install air conditioning in delivery vans under teamsters deal, 2023.
[396] Reuters. Ups, teamsters reach tentative deal, averting strike, 2023.
[397] Reuters. Ups workers vote to accept new contract including air conditioning, 2023.
[398] Reuters. Acropolis closes again as greek heatwave intensifies; tourists seek shade, 2024.
[399] Reuters. Greece closes more ancient tourist sites as heatwave persists, 2024.
[400] Reuters. Greece shuts acropolis as temperatures soar in early summer heatwave, 2024.
[401] Reuters. Iberians hit the beaches as temperatures rise 10°c above normal, 2024.
[402] Reuters. Indians battle intense heat with 'mad rush' for air conditioners and beer, 2024.
[403] Reuters. Mexico heat wave melts temperature records; cdmx among affected, 2024.
[404] Reuters. Mexico heat wave triggers exceptional power outages, 2024.
[405] Reuters. Scorching heat and lightning grip u.s. open second round, 2024.
[406] Reuters. Spain warns workers of heat risks as summer's first heatwave looms, 2024.
[407] Reuters. Water breaks, fans and ice as paris heatwave forces tweaks at the games, 2024.
[408] Reuters. France shuts schools as heatwave grips europe; sea off spain hits record high, 2025.
[409] Reuters. Greece shuts acropolis as heat soars, bans outdoor work in parts of country, 2025.
[410] Reuters. Marine heatwave pushes mediterranean temperatures up; european authorities restrict outdoor work and services, 2025.
[411] Reuters. Rhine shipping faces delays and surcharges as water levels drop in heat, 2025.
[412] Reuters. Rhine water levels fall again, forcing load restrictions on barges, 2025.
[413] Reuters. Scorching club world cup raises concerns for 2026; fifpro urges stronger heat protections, 2025.
[414] S. J. Rey, L. Anselin, P. Amaral, D. Arribas-Bel, R. X. Cortes, J. D. Gaboardi, W. Kang, E. Knaap, Z. Li, S. Lumnitz, et al. The pysal ecosystem: Philosophy and implementation. *Geographical Analysis*, 54(3):467–487, 2022.
[415] F. L. Ribeiro, J. Meirelles, F. F. Ferreira, and C. R. Neto. A model of urban scaling laws based on distance dependent interactions. *Royal Society open science*, 4(3):160926, 2017.
[416] F. L. Ribeiro and D. Rybski. Mathematical models to explain the origin of urban scaling laws. *Physics Reports*, 1012:1–39, 2023.
[417] J. Robert E. Lucas. On the mechanics of economic development. *Journal of Monetary*

Economics, 22(1):3–42, 1988.

[418] M. P. Roche. Taking innovation to the streets: microgeography, physical structure, and innovation. *Review of Economics and Statistics*, 102(5):912–928, 2020.

[419] J. Rokem and L. Vaughan. Segregation, mobility and encounters in jerusalem: The role of public transport infrastructure in connecting the 'divided city'. *Urban Studies*, 55(15):3454–3473, 2018.

[420] J. Rokem and L. Vaughan. Geographies of ethnic segregation in stockholm: The role of mobility and co-presence in shaping the 'diverse'city. *Urban Studies*, 56(12):2426–2446, 2019.

[421] M. Romanello, A. McGushin, C. Di Napoli, P. Drummond, N. Hughes, L. Jamart, H. Kennard, P. Lampard, B. S. Rodriguez, N. Arnell, et al. The 2021 report of the lancet countdown on health and climate change: code red for a healthy future. *The Lancet*, 398(10311):1619–1662, 2021.

[422] P. M. Romer. Endogenous technological change. *Journal of Political Economy*, 98(5):S71–S102, 1990.

[423] C. Roth, S. M. Kang, M. Batty, and M. Barthélemy. Structure of urban movements: polycentric activity and entangled hierarchical flows. *PloS one*, 6(1):e15923, 2011.

[424] R. Rothstein. *The color of law: A forgotten history of how our government segregated America*. Liveright Publishing, 2017.

[425] J. D. Sachs and A. M. Warner. The curse of natural resources. *European Economic Review*, 45(4-6):827–838, 2001.

[426] SafeGraph. "determining points of interest visits from location data: A technical guide to visit attribution". `https://www.safegraph.com/guides/visit-attribution-white-paper`.

[427] SafeGraph. Spend summary statistics. `https://docs.safegraph.com/docs/spend-summary-statistics`, 2024. Accessed: September 25, 2024.

[428] A. Sahasranaman and L. M. A. Bettencourt. Urban geography and scaling of contemporary Indian cities. *Journal of The Royal Society Interface*, 16(152):20180758, Mar. 2019. Publisher: Royal Society.

[429] A. Saiz. The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296, 2010.

[430] R. J. Sampson. *Great American City: Chicago and the Enduring Neighborhood Effect*. University of Chicago Press, 2012.

[431] R. J. Sampson and B. L. Levy. Beyond residential segregation: Mobility-based connectedness and rates of violence in large cities. *Race and social problems*, 12(1):77–86, 2020.

[432] M. Santamouris. Cooling cities: A review of heat mitigation techniques and their impacts. *Nature Reviews Earth & Environment*, 1:1–16, 2020. Review.

[433] C. Santana, F. Botta, H. Barbosa, F. Privitera, R. Menezes, and R. Di Clemente. Changes in the time-space dimension of human mobility during the covid-19 pandemic. *arXiv preprint arXiv:2201.06527*, 2022.

[434] A. Saxenian. *Regional advantage: Culture and competition in silicon valley and route 128, with a new preface by the author*. Harvard University Press, 1996.

[435] A. Schafer. The global demand for motorized mobility. *Transportation Research Part A: Policy and Practice*, 32(6):455–477, 1998.

[436] A. Schaffer, D. Muscatello, R. Broome, S. Corbett, and W. Smith. Emergency department visits, ambulance calls, and mortality associated with an exceptional heat wave in sydney, australia, 2011: a time-series analysis. *Environmental Health*, 11:1–8, 2012.

[437] T. C. Schelling. Bargaining, communication, and limited war. *Conflict Resolution*, 1(1):19–36, 1957.

[438] T. C. Schelling. Dynamic models of segregation. *Journal of mathematical sociology*, 1(2):143–186, 1971.

[439] T. C. Schelling. Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2):143–186, July 1971. Publisher: Routledge _eprint: https://doi.org/10.1080/0022250X.1971.9989794.

[440] M. Schiavina, S. Freire, A. Carioli, and K. MacManus. Ghs-pop r2023a - ghs population grid multitemporal (1975-2030), 2023.

[441] M. Schiavina, A. Moreno-Monroy, L. Maffenini, and P. Veneri. Ghs-fua r2019a - ghs

functional urban areas, derived from ghs-ucdb r2019a (2015), 2023.
[442] N. Schiff. Cities and product variety: evidence from restaurants. *Journal of Economic Geography*, 15(6):1085–1123, 2015.
[443] M. Schläpfer, L. M. Bettencourt, S. Grauwin, M. Raschke, R. Claxton, Z. Smoreda, G. B. West, and C. Ratti. The scaling of human interactions with city size. *Journal of the Royal Society Interface*, 11(98):20130789, 2014.
[444] M. Schläpfer, L. Dong, K. O'Keeffe, P. Santi, M. Szell, H. Salat, S. Anklesaria, M. Vazifeh, C. Ratti, and G. B. West. The universal visitation law of human mobility. *Nature*, 593(7860):522–527, 2021.
[445] W. Schlenker and M. J. Roberts. Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences*, 106(37):15594–15598, 2009.
[446] I. Schnell and N. Haj-Yahya. Arab integration in jewish-israeli social space: does commuting make a difference? *Urban Geography*, 35(7):1084–1104, 2014.
[447] A. Schrager. The introverts have taken over the us economy, 01 2024.
[448] J. A. Schumpeter. *Capitalism, Socialism and Democracy*. Harper, 1942.
[449] Secretaría de Estado de Transportes y Movilidad Sostenible. Estudio de movilidad de viajeros de ámbito nacional aplicando la tecnología Big Data: Informe metodológico. Technical report, Ministerio de Transportes y Movilidad Sostenible, 9 2024.
[450] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
[451] C. E. Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
[452] S. Silm and R. Ahas. Ethnic differences in activity spaces: A study of out-of-home nonemployment activities with mobile phone data. *Annals of the Association of American Geographers*, 104(3):542–559, 2014.
[453] S. Silm and R. Ahas. The temporal variation of ethnic segregation in a city: Evidence from a mobile phone use dataset. *Social Science Research*, 47:30–43, 2014.
[454] F. Simini, M. C. González, A. Maritan, and A.-L. Barabási. A universal model for mobility and migration patterns. *Nature*, 484(7392):96–100, 2012.
[455] F. Simini, A. Maritan, and Z. Néda. Human mobility in a continuum approach. *PLOS ONE*, 8(3):e60069, 2013.
[456] M. L. Small. *Unanticipated gains: Origins of network inequality in everyday life*. Oxford University Press, 2009.
[457] M. L. Small and L. Adler. The role of space in the formation of social ties. *Annual Review of Sociology*, 45(1):111–132, 2019.
[458] M. E. Smith. *Urban life in the distant past: The prehistory of energized crowding*. Cambridge University Press, 2023.
[459] R. M. Solow. A contribution to the theory of economic growth. *Quarterly Journal of Economics*, 70(1):65–94, 1956.
[460] E. Somanathan, R. Somanathan, A. Sudarshan, and M. Tewari. The impact of temperature on productivity and labor supply: Evidence from indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827, 2021.
[461] E. Somanathan, R. Somanathan, A. Sudarshan, and M. Tewari. The impact of temperature on productivity and labor supply: Evidence from indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827, 2021.
[462] C. Song, T. Koren, P. Wang, and A.-L. Barabási. Modeling the scaling properties of human mobility. *Nature Physics*, 6(10):818–823, 2010.
[463] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.
[464] M. Sonksen. The history of south central los angeles and its struggle with gentrification. *PBS SoCal (City Rising series)*.
[465] R. Squire. "what about bias in safegraph data?". `https://medium.com/safegraph/what-about-bias-in-the-safegraph-dataset-b95b7208893c`.
[466] A. J. Stier, M. G. Berman, and L. M. Bettencourt. Early pandemic covid-19 case growth rates increase with city size. *npj Urban Sustainability*, 1(1):31, 2021.
[467] A. J. Stier, K. E. Schertz, N. W. Rim, C. Cardenas-Iniguez, B. B. Lahey, L. M. Bettencourt,

and M. G. Berman. Evidence and theory for lower rates of depression in larger us urban areas. *Proceedings of the National Academy of Sciences*, 118(31):e2022472118, 2021.

[468] M. Storper. *The Keys to the City: How Economics, Institutions, Social Interaction, and Politics Shape Development*. Princeton University Press, 2013.

[469] S. A. Stouffer. Intervening opportunities: A theory relating mobility and distance. *American Sociological Review*, 5(6):845–867, 1940.

[470] A. Stutzer and B. S. Frey. Stress that doesn't pay: The commuting paradox. *Scandinavian Journal of Economics*, 110(2):339–366, 2008.

[471] M. Suárez, M. Delgado, and A. Murata. Why do the poor travel less? urban structure, economic informality and mobility in mexico city. *Urban Studies*, 53(12):2548–2569, 2016.

[472] Subdirección General de Planificación, Red Transeuropea y Logística. Controles de calidad realizados a los análisis de movilidad con BigData. Parte 1. Estudios básicos. Technical report, Ministerio de Transportes y Movilidad Sostenible, 10 2024.

[473] A. P. Susaneck. Boston: Roxbury (segregation by design). TU Delft Centre for the Just City, 2024.

[474] A. P. Susaneck. The bronx: The cross bronx expressway (segregation by design). TU Delft Centre for the Just City, 2024.

[475] Y. Tan, Y. Chai, and Z. Chen. Social-contextual exposure of ethnic groups in urban china: From residential place to activity space. *Population, Space and Place*, 25(7):e2248, 2019.

[476] R. Teskey, T. Wertin, I. Bauweraerts, M. Ameye, M. A. McGuire, and K. Steppe. Responses of tree species to heat waves and extreme heat events. *Plant, Cell & Environment*, 38:1699–1712, 2015.

[477] The Chicago Reporter. Chicago's 250 year history of segregation.

[478] The Economist. Global cities house-price index. `https://www.economist.com/graphic-detail/2019/03/11/global-cities-house-price-index`, 2019.

[479] The Economist. Global cities house-price index: Our interactive guide to housing data in the world's most desirable cities, Mar. 2019.

[480] B. Thrasher, W. Wang, A. Michaelis, F. Melton, T. Lee, and R. Nemani. Nasa global daily downscaled projections, cmip6. *Scientific data*, 9(1):262, 2022.

[481] E. Times. Dmrc ensures cool commutes at 24°c; daily ridership hits 60.17 lakh in may, 2024.

[482] F. Times. Can india adapt to extreme heat? FT Film video, 2022. Directed and produced by Juliet Riddell and Sidrah Fatma Ahmed; includes first-hand interviews with farmers, students, business owners, factory workers and scientists in India on how rising temperatures affect working conditions and productivity.

[483] F. Times. Gig workers in dubai get access to cooling d̈eliveroo busämid heat stress. *Financial Times*, 2024. Picture of Deliveroo bus as cool refuge for gig workers.

[484] F. Times. Glovo pauses 'heat bonuses' for italian delivery workers after backlash, 2025.

[485] H. Times. Heatstroke alert in delhi: what to do (avoid mid-day exertion; stay indoors), 2024.

[486] H. Times. Working in extreme heat: Delhi's auto-rickshaw drivers improvise, take breaks, 2024.

[487] A. Tobías, M. Hashizume, Y. Honda, F. Sera, C. F. S. Ng, Y. Kim, D. Roye, Y. Chung, T. N. Dang, H. Kim, et al. Geographical variations of the minimum mortality temperature at a global scale: a multicountry study. *Environmental epidemiology*, 5(5):e169, 2021.

[488] I. Today. Heatwave burns hole in delhi street vendors' wallets: survey, 2024.

[489] TomTom. How have traffic patterns changed across the world as a result of covid-19?, 2020. Accessed 2025-09-23.

[490] G. Tóth, J. Wachs, R. Di Clemente, Á. Jakobi, B. Ságvári, J. Kertész, and B. Lengyel. Inequality is rising where social network segregation interacts with urban topology. *Nature communications*, 12(1):1143, 2021.

[491] Transformative Urban Mobility Initiative (TUMI). E-bus mission city factsheets: Delhi, mumbai, bangalore (india); mexico city, guadalajara, monterrey (mexico); jakarta (indonesia). `https://transformative-mobility.org/`, 2022. Multiple city factsheets (PDFs), 2022–2024, produced by GIZ/TUMI. Used for city-level modal split shares. Accessed 28 September 2025.

[492] A. Tversky and D. Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
[493] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981.
[494] UAE Government. Mid-day break rule for workers under the sun, 2025.
[495] D. o. E. United Nations and S. Affairs. World urbanization prospects: The 2018 revision. `https://population.un.org/wup/`, 2019. p. iv: data consistent with medium variant only.
[496] D. o. E. United Nations and S. Affairs. World population prospects 2024 revision—high variant. `https://population.un.org/wpp/`, 2024. Accessed August 2025.
[497] U.S. Bureau of Labor Statistics. Databases, tables calculators by subject. `https://www.bls.gov/data/`, 2022. Metropolitan GMP data, available from the BLS Data Portal.
[498] U.S. Census Bureau. American Community Survey. `https://www.census.gov/programs-surveys/acs/`, 2024. Accessed: 2023-05-16.
[499] U.S. Census Bureau. Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES) Data. `https://lehd.ces.census.gov/data/`, 2024. Accessed: 2023-05-16.
[500] S. Van Nieuwerburgh. The remote work revolution: Impact on real estate values and the urban environment: 2023 areuea presidential address. *Real Estate Economics*, 51(1):7–48, 2023.
[501] N. T. Vargas, Z. J. Schlader, O. Jay, and A. Hunter. Prioritize research on human behaviour during extreme heat. *Nature Human Behaviour*, 7(4):473–474, 2023.
[502] vCity. Climate shelters in barcelona: Accessibility for vulnerable populations, 2024. Part of the Proximity City Living Lab within the vCity project.
[503] A. M. Vicedo-Cabrera, N. Scovronick, F. Sera, D. Royé, R. Schneider, A. Tobias, C. Astrom, Y. Guo, Y. Honda, D. Hondula, et al. The burden of heat-related mortality attributable to recent human-induced climate change. *Nature climate change*, 11(6):492–500, 2021.
[504] F. Wang and C. Chen. On data processing required to derive mobility patterns from passively-generated mobile phone data. *Transportation Research Part C: Emerging Technologies*, 87:58–74, 2018.
[505] Q. Wang, N. E. Phillips, M. L. Small, and R. J. Sampson. Urban mobility and neighborhood isolation in america's 50 largest cities. *Proceedings of the National Academy of Sciences*, 115(30):7735–7740, 2018.
[506] S. Wang, Y. Zheng, G. Wang, T. Yabe, E. Moro, and A. entland. Infrequent activities predict economic outcomes in major american cities. *Nature Cities*, 1(4):305–314, 2024.
[507] X. Wang, R. Xue, M. Lu, and J. Wu. Riders under the heat: Exploring the impact of extreme heat on the integration of bike-sharing and public transportation in shenzhen, china. *ISPRS International Journal of Geo-Information*, 13(12):438, 2024.
[508] T. Watson. Inequality and the measurement of residential segregation by income in american neighborhoods. *Review of Income and Wealth*, 55(3):820–844, 2009.
[509] A. Wesolowski, N. Eagle, A. M. Noor, R. W. Snow, and C. O. Buckee. The impact of biases in mobile phone ownership on estimates of human mobility. *Journal of the Royal Society Interface*, 10(81):20120986, 2013.
[510] L. Whitmarsh, K. Mitev, N. Nash, C. Hoolohan, W. Poortinga, C. Whittle, P. Haggar, C. Tipping, H. Dudley, L. Turner, et al. "moments of change" and low-carbon behaviors: A multidisciplinary, systematic review. *Wiley Interdisciplinary Reviews: Climate Change*, 16(4):e70014, 2025.
[511] B. Wilson, C. Neale, and J. Roe. Urban green space access, social cohesion, and mental health outcomes before and during covid-19. *Cities*, 152:105173, 2024.
[512] W. J. Wilson. *The truly disadvantaged: The inner city, the underclass, and public policy*. University of Chicago Press, 2012.
[513] D. W. Wong and S.-L. Shaw. Measuring segregation: An activity space approach. *Journal of geographical systems*, 13(2):127–145, 2011.
[514] S. N. Wood. mgcv: Gams and generalized ridge regression for r. *R news*, 1(2):20–25, 2001.
[515] World Bank. Climate change knowledge portal. `https://climateknowledgeportal.`

worldbank.org, 2024. Accessed September 30, 2024.

[516] Y. Wu, B. Wen, A. Gasparrini, B. Armstrong, F. Sera, E. Lavigne, S. Li, Y. Guo, A. Overcenco, A. Urban, et al. Temperature frequency and mortality: Assessing adaptation to local temperature. *Environment international*, 187:108691, 2024.

[517] H. Xiao, A. Wu, and J. Kim. Commuting and innovation: Are closer inventors more productive? *Journal of Urban Economics*, 121:103300, 2021.

[518] C. Xiong, S. Hu, M. Yang, W. Luo, and L. Zhang. Mobile device data reveal the dynamics in a positive relationship between human mobility and covid-19 infections. *Proceedings of the National Academy of Sciences*, 117(44):27087–27089, 2020.

[519] Y. Xu, A. Belyi, P. Santi, and C. Ratti. Quantifying segregation in an integrated urban physical-social space. *Journal of The Royal Society Interface*, 16(160):20190536, Nov. 2019. Publisher: Royal Society.

[520] Y. Xu, L. E. Olmos, D. Mateo, A. Hernando, X. Yang, and M. C. González. Urban dynamics through the lens of human mobility. *Nature computational science*, 3(7):611–620, 2023.

[521] Z. Xu, R. A. Etzel, H. Su, C. Huang, Y. Guo, and S. Tong. The impact of heat waves on children's health: a systematic review. *International Journal of Biometeorology*, 58:239–247, 2014.

[522] T. Yabe, B. G. B. Bueno, X. Dong, A. Pentland, and E. Moro. Behavioral changes during the covid-19 pandemic decreased income diversity of urban encounters. *Nature communications*, 14(1):2310, 2023.

[523] T. Yabe, B. García Bulle Bueno, M. R. Frank, A. Pentland, and E. Moro. Behaviour-based dependency networks between places shape urban economic resilience. *Nature human behaviour*, pages 1–11, 2024.

[524] X.-Y. Yan, C.-X. Zhao, Y. Fan, Z. Di, and W.-X. Wang. Universal predictability of mobility patterns in cities. *Journal of the Royal Society Interface*, 11(98):20140834, 2014.

[525] L. Yang, D. Holtz, S. Jaffe, S. Suri, S. Sinha, J. Weston, C. Joyce, N. Shah, K. Sherman, B. Hecht, and J. Teevan. The effects of remote work on collaboration among information workers. *Nature Human Behaviour*, 6(1):43–54, 2022.

[526] Q. Yin, J. Wang, Z. Ren, J. Li, and Y. Guo. Mapping the increased minimum mortality temperatures in the context of global climate change. *Nature communications*, 10(1):4640, 2019.

[527] W. You. The economics of speed: the electrification of the streetcar system and the decline of mom-and-pop stores in boston, 1885–1905. *American Economic Journal: Applied Economics*, 13(4):285–324, 2021.

[528] Y. Zahavi and A. Talvitie. *Regularities in travel time and money expenditures*. Number 750. 1980.

[529] K. K. Zander, W. J. Botzen, E. Oppermann, T. Kjellstrom, and S. T. Garnett. Heat stress causes substantial labour productivity loss in australia. *Nature climate change*, 5(7):647–651, 2015.

[530] A. Zanobetti, M. S. O'neill, C. J. Gronlund, and J. D. Schwartz. Summer temperature variability and long-term survival among elderly people with chronic disease. *Proceedings of the National Academy of Sciences*, 109(17):6608–6613, 2012.

[531] K. Zhang, Y. Chen, and C. E. Begley. Impact of the 2011 heat wave on mortality and emergency department visits in houston, texas. *Environmental Health*, 14:1–7, 2015.

[532] P. Zhang, O. Deschenes, K. Meng, and J. Zhang. Temperature effects on productivity and factor reallocation: Evidence from a half million chinese manufacturing plants. *Journal of Environmental Economics and Management*, 88:1–17, 2018.

[533] W. Zhang, M. N. del Prado, V. Gauthier, and S. Milusheva. The netmob2024 dataset: Population density and od matrices from four lmic countries. *arXiv preprint arXiv:2410.00453*, 2024.

[534] Q. Zhao, S. Li, T. Ye, Y. Wu, A. Gasparrini, S. Tong, A. Urban, A. M. Vicedo-Cabrera, A. Tobias, B. Armstrong, et al. Global, regional, and national burden of heatwave-related mortality from 1990 to 2019: A three-stage modelling study. *PLoS medicine*, 21(5):e1004364, 2024.

[535] L. Zhong, L. Dong, Q. R. Wang, C. Song, and J. Gao. Universal expansion of human mobility across urban scales. *Nature Cities*, 2:603–607, 2025.

[536] C. D. Ziter, E. J. Pedersen, C. J. Kucharik, and M. G. Turner. Scale-dependent interactions between tree canopy cover and impervious surfaces reduce daytime urban heat during summer. *PNAS*, 116(15):7575–7580, 2019.

[537] P. Zschenderlein, A. H. Fink, S. Pfahl, and H. Wernli. Processes determining heat waves across different european climates. *Quarterly Journal of the Royal Meteorological Society*, 145(724):2973–2989, 2019.