Year : 2023

THREE ESSAYS IN URBAN AND DEVELOPMENT ECONOMICS USING SPATIAL MICRO-DATA

Nöbauer Bernhard Florian

Nöbauer Bernhard Florian, 2023, THREE ESSAYS IN URBAN AND DEVELOPMENT ECONOMICS USING SPATIAL MICRO-DATA

Originally published at : Thesis, University of Lausanne

Posted at the University of Lausanne Open Archive http://serval.unil.ch
Document URN : urn:nbn:ch:serval-BIB_B13A289F92352

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

THREE ESSAYS IN URBAN AND DEVELOPMENT ECONOMICS USING SPATIAL MICRO-DATA

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Doctorat ès Sciences Économiques,
mention « Économie politique »

par

Bernhard Florian NÖBAUER

Directeur de thèse
Prof. Marius BRÜLHART

Jury

Prof. Boris NIKOLOV, président
Prof. Dzhamilya NIGMATULINA, experte interne
Prof. Gilles DURANTON, expert externe

LAUSANNE
2023

UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMIE

THREE ESSAYS IN URBAN AND DEVELOPMENT ECONOMICS USING SPATIAL MICRO-DATA

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Études Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Doctorat ès Sciences Économiques,
mention « Économie politique »

par

Bernhard Florian NÖBAUER

Directeur de thèse
Prof. Marius BRÜLHART

Jury

Prof. Boris NIKOLOV, président
Prof. Dzhamilya NIGMATULINA, experte interne
Prof. Gilles DURANTON, expert externe

LAUSANNE
2023

UNIL | Université de Lausanne
Faculté des hautes études commerciales
Internef - 261
CH-1015 Lausanne

IMPRIMATUR

La Faculté des hautes études commerciales de l'Université de Lausanne autorise l'impression de la thèse de doctorat rédigée par

Bernhard Nöbauer

intitulée

Three Essays in Urban and Development Economics Using Spatial Micro-Data

sans se prononcer sur les opinions exprimées dans cette thèse.

Lausanne, le 21.09.2023

Professeure Marianne Schmid Mast, Doyenne

Thesis Committee

Prof. Marius BRÜLHART

Université de Lausanne

Thesis supervisor

Prof. Gilles DURANTON

Wharton School of the University of Pennsylvania

External member of the doctoral committee

Prof. Dzhamilya NIGMATULINA

Université de Lausanne

Internal member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Bernhard Florian NÖBAUER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature:  Date: 3 July 2023

Prof. Marius BRÜLHART
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Bernhard Florian NÖBAUER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature:  _____ Date: __1st July 2023__

Prof. Gilles DURANTON
External expert

University of Lausanne
Faculty of Business and Economics

PhD in Economics
Subject area Political Economy

I hereby certify that I have examined the doctoral thesis of

Bernhard Florian NÖBAUER

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature: Dzhamilya Nigmatulina Date: 3rd of July 2023

Prof. Dzhamilya NIGMATULINA
Internal expert

Spatial Data in Economics: Three Essays

Bernhard Nöbauer
Unversité de Lausanne

September 2023

Acknowledgements

This thesis would not have been possible without the many people who accompanied me on this long and sometimes difficult path.

First and foremost, I want to express my deep gratitude to my supervisor, Marius Brülhart. Thank you for your guidance, countless hours of discussions, and innumerable helpful comments. You might humbly suggest that you are just doing your job, but you are doing much more than that. Thank you for being the encouraging, constructive, and kind person you are. It was a pleasure to be your student.

I also want to thank Gilles Duranton, who hosted me for six months at the Wharton School of the University of Pennsylvania and whose work inspired me throughout my PhD. You had already tackled virtually every challenge I encountered at some point in your distinguished career, and you still took more time than I could have hoped for to discuss my work with all its facets in depth.

Jamila Nigmatulina completed my doctoral committee and offered insightful questions at the defense and many helpful comments before and after. I further want to thank Mélise Jaud for allowing me to gain fascinating insights into the situation in Myanmar and valuable experiences during my stint as a short-term consultant at the World Bank.

The economics department at the Université de Lausanne offered great conditions to pursue my thesis. I was able to share this experience with many amazing colleagues. I also had a fantastic cohort during the coursework at the Study Center Gerzensee, making me feel welcome from the beginning of the PhD. I am especially thankful to my long-term office mates Tobias Lehmann and Jeremias Kläui for many fruitful discussions and their friendship.

I also want to thank my parents Eva and Rupert. Thank you for your unconditional support over my whole educational career. Finally, I want to thank my wonderful partner, Nora. You were a constant cause for joy over all this time. Thank you all for your love.

Contents

Introduction

Economic models often do not explicitly include a spatial dimension. However, economic activity typically happens at a clearly defined location, and in many circumstances, this location matters. Consider the example of housing, on which households spend more than on any other expenditure item.[1] Two identical apartments can be priced very differently if one is in a sought-after neighborhood in a thriving metropolis while the other is at the periphery of a town with a declining workforce. Moving from one location to another affects the available jobs, leisure activities, goods, and services. Certain places have many economic opportunities, others have a high crime level or conflict - and sometimes both exist at the same location.

Over the last decade, more data have been collected than ever before (Reinsel et al., 2018). To the extent that these data can be accessed, they open exciting possibilities for academic research. A considerable share of the data available today have a geographical dimension, allowing us to work on topics like the ones mentioned above. Some data can be matched to low-level administrative units, while others are associated with precise geographical coordinates. In the past, spatially disaggregated data have been primarily available for individual high-income countries. Today, there is an increasing amount of such data in the context of developing countries. Moreover, some data cover multiple countries or even the entire world. These data allow for international comparisons that are less affected by data comparability across countries. In my PhD thesis, I use these advancements in data availability to revisit longstanding questions in empirical microeconomics. I write about topics that can be attributed to urban economics, development economics, and the economics of conflict.

Almost all data sources I use for my work are publicly available and free to use. The only exception is a dataset on short-term rental objects from Airbnb that I obtained from AirDNA, a company specialized in "short-term rental data and analytics ".[2] This dataset allows me to use information about the price, size, amenities, and geographic coordinates of over three million properties worldwide. The information is available in an internationally standardized way, with a property in Chile being described with the same variables as a property in Nigeria or Indonesia. In the literature, housing cost data are usually taken from national statistical offices or from country-specific real estate platforms[3] and using these novel data allows me to go beyond that.

[1] `https://ec.europa.eu/eurostat/web/products-eurostat-news/-/DDN-20181130-1`, last accessed: 2023-08-28.
[2] `https://airdna.co`, last accessed: 2023-01-16.
[3] Examples of such platforms include `https://www.zillow.com` for the United States, `https://www.seloger.com` for

I use the data to research regularities of price levels i) within cities, as a function of distance to the city center, and ii) between cities. Thereby, prices of Airbnb properties can serve as a proxy for housing costs more broadly, as people offering short-term and long-term rental objects compete for the same properties. France and the United States provide long-term rental data on a fine geographical scale, and I show that these rents are correlated with the short-term rents from Airbnb.

Different countries use different city definitions. However, since I am working with a worldwide sample, it is advantageous to define city boundaries and city centers in an internationally consistent way. In that regard, too, I benefit from the large amount of data that is publicly available today. My definition starts with so-called city tags from OpenStreetMap. I spatially join them to the urban center database from the global human settlement project of the European Commission. After applying different cutoffs (e.g., requiring a city to have at least 300,000 inhabitants and 100 Airbnb properties), I am left with a sample of 734 cities in 123 countries.

The city tags in OpenStreetMap are based on crowd intelligence, with users being asked to place them "at the center of the city, like the central square, a central administrative or religious building or a central road junction".[4] This makes the tags promising candidates for the center of a city. Together with a research assistant, I manually evaluated all of them, complementing them with the centers proposed by Google Maps. We also provide our own best guess whenever the proposed center does not seem to be a good choice. The literature offers substantial work concerning the delineation of cities, but research on the locations of city centers is very scarce, especially on a worldwide scope.

In *The monocentric city model worldwide: Rent, density, and transportation cost gradients in 734 cities*, I use these data and definitions to estimate urban rent gradients for an international cross-section of cities. To do so, I estimate the effect of distance to the city center on housing prices using a log-log specification. I find this specification to be less affected by city delineations than the log-linear specification that is also often used in the literature. A potential problem arises as properties close to the city center can be quite different from properties at a city's periphery. To alleviate this, the prices I use for the estimation are residuals from a hedonic regression.

I find an average rent gradient of -0.064, implying a 0.64% increase in price for every 10% increase in distance. Rent gradients are less pronounced in cities that are smaller and in cities located by a large water body. Moreover, there is a non-linear relation with income, with upper-middle-income countries exhibiting the flattest average rent gradients. I also estimate population density gradients based on the same city definitions and on population data from the Global Human Settlement project. Density gradients are steeper than rent gradients for most cities, with an average value of -0.356.

France, or `https://www.homegate.ch` for Switzerland (all last accessed: 2023-08-28).

[4] `https://wiki.openstreetmap.org/wiki/Tag:place%3Dcity`, last accessed: 2023-01-13.

Using the monocentric city model, I then use the rent gradients and the density gradients to infer information about the elasticity of intra-city transportation costs with respect to distance. To do so, I use a standard version of the model and amend it with the common assumption that rents and population density both follow a log-log relation with distance. I show that this setup results in the transportation cost gradient being equal to the difference between the rent gradient and the density gradient.

I find an average elasticity of transportation costs to distance from the city center of 0.3. This aligns with a concave transportation cost function, while the common simplifying assumption of linear transportation costs would imply an elasticity of 1. However, the average value masks considerable heterogeneity. For example, the average among French cities is 0.47, while the average among cities in the United States is 0.07. The latter precisely matches the urban cost estimate that Duranton and Puga (2022) obtain when performing a similar exercise for cities in the United States. However, my global estimate of 0.3 suggests that the US is an outlier, with transportation costs being less sensitive to distance than elsewhere.

After considering differences in housing costs within cities, I turn to analyze differences in housing costs between cities in *How much more expensive is housing in larger cities? Worldwide evidence from Airbnb*. The literature proposes extensive evidence on agglomeration economies, reporting higher productivity and wages in larger cities. However, since we have not converged to live in a single gigantic city, there must be disadvantages, with housing costs being a prominent example. In this paper, I measure the elasticity of housing costs with respect to city size on the same worldwide sample.

Methodologically, I follow Combes et al. (2018) in measuring housing costs at the city center, which has the advantage that differences in transportation costs do not significantly influence comparisons across cities. One could also compare average housing costs in different cities. However, the average property in a big city like Tokyo is much further away from the center than in a smaller city like Kagoshima. Therefore, its inhabitants face higher transportation costs that would confound the comparison.

Using a hedonic regression with city fixed effects and city-specific distance gradients, I create a ranking of the 734 cities regarding the rental rate of a representative property at the city center. Given my methodology and data, I estimate Amsterdam, London, New York, and San Francisco to have the highest rental prices in the world. Caracas, Mandalay, Monteria, and Srinagar are at the other end of the ranking, with rental prices that are around 20 times lower.

In the second stage, I regress the log of the estimated rental prices on log population while controlling for log area. I include country-fixed effects, so the coefficients are estimated from within-country variation. Moreover, I control for various city characteristics, including the number of Airbnb properties per 100,000 inhabitants, to control for the attractiveness of a city to tourists. An instrumental variable approach in which historical population sizes are used as an instrument serves as a robust-

ness check. In my preferred specification, I estimate an elasticity of 0.161. This coefficient implies that a 10% higher population size is associated with housing costs that are 1.61% higher. The effect is statistically significant at the 1% level.

This worldwide estimate is in the ballpark of the few estimates the literature provides for individual countries. However, once again, there is considerable heterogeneity. The estimated elasticity is above the global average for the United States and Russia and is particularly high in the eurozone and India. Stringent planning regulations might be an explanation for this, limiting the vertical and horizontal margins a city has available to accommodate an influx of people. The Rosen-Roback model (Rosen, 1979; Roback, 1982) offers an additional explanation: The literature reports stronger agglomeration economies in developing countries. Taken together with my finding that high-income countries tend to exhibit rents that increase faster in city size, this suggests that high-income countries have amenities that increase faster (or decrease slower) in city size than in other countries. However, being surrounded by many people might not always be beneficial. I find that denser cities are cheaper in the Mexican context. I offer suggestive evidence that crime might explain this unusual result.

In the two papers discussed above, novel property-level data allow me to go beyond one country and compare cities worldwide. However, better data can also help to zoom in on a particular setting. In early studies, the relationship between exports of natural resources and conflict was typically assessed with entire countries as the unit of observation. More detailed data allow researchers to work with subnational geographical entities as the unit of observation. Such within-country studies have the advantage that the political and institutional setting is more comparable between observations.

With *Trade and conflict in Myanmar: A reverse China shock*, I contribute to this literature. I zoom in on Kachin and Shan, the two states of Myanmar that border China. These states feature many metal mines and suffer from a high level of conflict. I use event-level conflict data from ACLED (Raleigh et al., 2010), VIIRS night light data (Elvidge et al., 2021), as well as a dataset on mining licenses provided by MCRB (2022) that has, to my knowledge, not yet been used for academic research. These datasets are either disaggregated at a low administrative level or geocoded, allowing me to work at the township level.

I look at Myanmar's exports of mining goods to China from 2012 to 2020. Over this time, 86% of the value obtained by exporting mining goods was due to trade with China. On the other hand, Myanmar accounted for only 0.2% of the value of Chinese imports in mining goods.[5] This imbalance allows me to use Chinese demand shocks to identify the effect of exports of mining goods on local conflict in the mines' townships in Kachin and Shan. Thereby, I essentially mirror the estimation strategy of Autor et al. (2013) and apply it to the setting of a developing country.

First, I disaggregate nationwide export values in different metals to the townships hosting mines

[5] These numbers come from BACI, as do the trade data I use for my calculations.

associated with these metals. Second, I construct an instrumental variable to deal with endogeneity concerns. In particular, there might be reverse causality, with conflict events negatively affecting exports. To construct the instrument, I use Chinese imports from other low and middle-income countries and interact them with the local mining geography in Kachin and Shan. Effects are assessed using regressions with township-fixed effects and year-fixed effects. Export exposure in mining goods increases the number of conflict events, with an elasticity of 0.46 for the OLS specification. The elasticity increases to around 0.54 when using the IV approach. Ethnicity seems to matter, as most of the increase in conflict events is due to townships not inhabited by the nationwide ethnic majority.

A placebo test further validates the results, replacing trade flows in mining goods with those in other goods of comparable importance for the export sector, typically produced elsewhere. As expected, the effect disappears when using this counterfactual measure of export exposure in placebo goods. Furthermore, I analyze the effect of export exposure on night lights as a proxy of economic development. In the immediate neighborhood of mines, pixels are brighter in years with higher export exposure. However, the effect does not spill over to nearby areas. Areas inhabited by ethnic minorities experience a lower increase in night lights when export exposure is higher, with the effect disappearing even quicker with distance.

All three papers in this thesis ask economic questions that have been around for a long time and that I readdressed with spatial data that became available relatively recently. While they cover different topics, they share the importance of geography. None of the questions would make sense in a void where people could move and trade without any transport costs. It is likely that the surge in data creation continues and that a considerable share of these data will have a spatial component. This development will create challenges concerning privacy protection and concerning the availability of these data for academic research and the general public. However, it will also offer exciting new opportunities to enhance our understanding of the economy and society more generally.

Bibliography

Autor, D.H., Dorn, D., and Hanson, G.H., 2013. The China syndrome: Local labor market effects of import competition in the United States, *American Economic Review*, 103 (6), 2121–2168.

Combes, P.P., Duranton, G., and Gobillon, L., 2018. The costs of agglomeration: House and land prices in French cities, *The Review of Economic Studies*, 86 (4), 1556–1589.

Elvidge, C.D., Zhizhin, M., Ghosh, T., Hsu, F.C., and Taneja, J., 2021. Annual time series of global viirs nighttime lights derived from monthly averages: 2012 to 2019, *Remote Sensing*, 13 (5), 922.

MCRB, 2022. MEITI licenses explorer for 6th MEITI summary data report, Online, Myanmar Centre for Responsible Business, `https://meiti-map.org` (last accessed: 07.03.2023).

Raleigh, C., Linke, A., Hegre, H., and Karlsen, J., 2010. Introducing ACLED: An armed conflict location and event dataset: Special data feature, *Journal of Peace Research*, 47 (5), 651–660.

Reinsel, D., Gantz, J., and Rydning, J., 2018. Data age 2025. The digitalization of the world, from edge to core, White Paper, International Data Corporation.

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

Rosen, S., 1979. Wage-based indexes of urban quality of life, *in:* P. Mieszkowski and M. Straszheim, eds., *Current Issues in Urban Economics*, John Hopkins University Press, 74–104.

Chapter 1

The monocentric city model worldwide: Rent, density, and transportation cost gradients in 734 cities

The monocentric city model worldwide: Rent, density, and transportation cost gradients in 734 cities[1]

Bernhard Nöbauer[2]
University of Lausanne

September 2023

Abstract

I use the monocentric city model to relate urban rent and density gradients to a third gradient describing transportation costs. I estimate rent gradients for 734 cities worldwide using internationally comparable data on Airbnb properties. The average elasticity of rent to distance from the city center is -0.06. Rent gradients are less pronounced in cities that are smaller or located in upper-middle-income countries. Density gradients are steeper than rent gradients in most cities. Combining the two types of gradients, I estimate the elasticity of transportation costs to distance from the city center to be 0.3 on average. Taken together, the two estimates imply a concave transportation cost function. While I precisely match the Duranton and Puga (2022) urban cost estimate of 0.07 for the United States, my global estimate of 0.3 suggests that the US is an outlier, with transportation costs being less sensitive to distance than elsewhere.

[1] I thank Marius Brülhart for his constant invaluable guidance and Gilles Duranton for his hospitality and many precious discussions. Moreover, I thank Prottoy Akbar, Sophie Calder-Wang, Dzhamilya Nigmatulina, Diego Puga, Jeff Wooldridge, my colleagues at the University of Lausanne, the participants of the Urban Lunch at the Wharton School of the University of Pennsylvania, and the participants of the 12th European Meeting of the Urban Economics Association for their helpful comments. Finally, I thank Laura Camarero Wislocka for her excellent help with assessing and defining city centers.

[2] Departement of Economics, Faculty of Business and Economics (HEC Lausanne), University Of Lausanne, 1015 Lausanne, Switzerland; `bernhard.noebauer@bluewin.ch`.

1 Introduction

The monocentric city model illustrates stylized facts about cities and provides intuition on how rents, transportation costs, and living space are interrelated. It has a long tradition, is typically the first model that is taught in university courses on urban economics, and still possesses importance within and beyond the field (Duranton and Puga, 2015). Nevertheless, the empirical evidence related to the monocentric city model is comparatively scarce. This paper helps fill this gap by estimating rent and density gradients for 734 cities worldwide. To do this, I use internationally consistent data, including data about short-term rental properties from Airbnb. The paper then uses the structure of the monocentric city model to relate these two gradients to infer information about the urban transportation cost gradient.

Transportation costs are one of the most relevant factors in the study of cities. The inherent advantage of cities is proximity. Being physically close to other people allows individuals to exchange knowledge and goods, sustain infrastructure, host events, and engage with one another. If traveling from A to B would cost neither time nor money, cities might not exist. Given that it is costly, people face a trade-off between being closer to advantageous locations versus having cheaper unit costs of housing. But how should we represent transportation costs in economic models?

A common simplifying assumption is to model transportation costs as being linear in distance, implying that the monetary and non-pecuniary costs of traveling 10 km are twice as high as those of traveling 5 km. One example of this are iceberg trade costs. In contrast, I come to the conclusion that most cities feature concave transportation costs, implying that travelling a distance twice as long is less than twice as costly. The average elasticity of transportation costs with regard to distance is around 0.3, while the linear case would suggest an elasticity of 1. This suggests that the monetary and non-pecuniary costs of traveling 10 km are in fact only 30% higher as those of traveling 5 km. Intuitive explanations for concave transportation costs include higher congestion and more frequent public transport stops in more central parts of a city, as well as fixed costs in order to get to a bus stop or into the car. While transportation costs are concave for the overwhelming majority of cities, the elasticity varies considerably across and within countries.

Instead of directly measuring transportation costs for one particular mode of transportation as e.g. in Akbar et al. (2021), I infer them from market prices and location choices that are structurally connected by the monocentric city model. I use a standard version of this model and amend it with the assumption that rents and population density both follow a log-log relationship with distance. I show empirically that this functional form is considerably more robust to the precise delineation of cities than a log-linear relationship, which is the other option commonly used in the corresponding empirical literature. This setup results in the transportation cost gradient being equal to the difference between the rent gradient and the density gradient. I estimate these two gradients for a worldwide cross-section of cities.

Duranton and Puga (2015) note that "the empirical knowledge accumulated on the monocentric urban model and its extensions remains limited" and that "[the] literature has often struggled

to provide evidence of negative gradients for the unit price of housing" (p. 523). The available studies are based on an individual city or one single country at most. Examples include Gupta et al. (2022) for the United States, Combes et al. (2018) for France, or Ding and Zhao (2014) and Li and Wan (2021) for Beijing. For a long time, there existed simply no internationally comparable data for rents or house prices on an object-level basis, as is needed for such an analysis.[1]

To my knowledge, the only other project that estimates rent gradients on a global scale is the recent work by Liotta et al. (2022), who use data from different real estate websites for 192 cities worldwide. They estimate the elasticity of rental prices (and population density) on income net of transportation costs. In comparison to their work, this paper relies on real estate data from a single internationally consistent source. Moreover, I use a first-stage hedonic regression to separate out the effect of various property-level characteristics. This is important, as properties close to the city center might be quite different from properties on the outskirts. If the former are more expensive, we want to know to which extent this is due to the more central location, and to which extent these properties have more rooms or nicer amenities. In contrast, if they are smaller and have less amenities, this might imply an underestimation of the effect of distance. Finally, Liotta et al. (2022) aggregate their rental data on the grid cell level, while I work directly with the coordinates of the objects.

I find an average rent gradient of -0.064, implying a 0.64% increase in price for every 10% increase in distance. Contrasting my estimates for different subsamples with the empirical literature mentioned above, I find them to be roughly comparable. Combes et al. (2018) report that house price gradients decrease in city size in France. I confirm this finding for most countries (that have at least 10 cities in my sample), with the notable exception of countries in Latin America. Moreover, rent gradients are flatter in cities in upper-middle-income countries and in cities situated near a major water body.

The monocentric city model predicts negative rent and density gradients. If these predictions hold, the density gradient needs to be steeper than the rent gradient to yield a sensible transportation cost function in my setting, i.e. one that features positive transportation costs that are increasing in distance. This is indeed the case for 89% of cities, with an average density gradient of -0.356. Combining the two gradients yields an average transportation cost gradient of about 0.3, implying decreasing marginal transportation costs.

This average value masks considerable heterogeneity. For example, the average among French cities is 0.47, while the average among cities in the United States is 0.07. The latter precisely matches the urban cost estimate that Duranton and Puga (2022) obtain when performing a similar exercise for cities in the United States, although they disregard population density and focus solely on price. While the similarity between our estimates is reassuring for the robustness of the result for the US, my findings suggest that the United States is an outlier. Internationally,

[1] Data on Airbnb properties have been used extensively to research the effect of short-term rental objects on the long-term rental market (see, for example, Calder-Wang, 2021; Àngel Garcia-López et al., 2020; Barron et al., 2021). Almagro and Domínguez-Iino (2022) use such data to examine endogenous amenities and their distributional effects. Coles et al. (2018) briefly discuss rent gradients in the context of New York City. However, these papers are all concerned with an individual city or, at most, different cities in an individual country.

the average elasticity of transportation costs to distance is about four times higher. In the context of their model, this would imply considerably higher urban costs. However, even within the US, the average estimate masks substantial heterogeneity, with the elasticity of transportation costs being much higher in cities like New York, Boston, or Chicago.

The remainder of the paper is organized as follows: Section 2 specifies the theory, while Section 3 defines the cities and their centers. Sections 4, 5, and 6 present the results on rent gradients, density gradients, and transportation cost gradients, respectively. Section 7 provides a discussion, and section 8 concludes. Finally, the appendix contains a more in-depth description of my city definitions and auxiliary results.

2 Data and city definitions

The data on Airbnbs come from AirDNA, a company specialized in "short-term rental data and analytics".[2] They contain close to all properties that were advertised on Airbnb at least once between 01.01.2018 and 25.03.2019, over 9.4 million properties in total.[3] By combining information about days for which properties are rented with information about the properties' prices for these days, AirDNA is able to estimate the prices actually paid by customers. For every property, I have information about the average daily price over the twelve months before the date on which a property was last web scraped from the Airbnb website. I also have the coordinates of the location for each property, even though some of them are scrambled by at most 500 meters due to security concerns. Moreover, the data contain a substantial number of covariates, from the number of bedrooms to the presence of a hairdryer. All of these variables are available in an internationally standardized way. After dropping properties that were never rented over the period in question, my dataset still contains more than 3 million properties within the cities defined below.[4]

It is not straightforward to determine where a city ends and where its center is located. Different countries have different rules to set administrative boundaries. Official city delineations are, therefore, hardly comparable. Fortunately, there have been attempts to find internationally consistent definitions of cities on which I can build. A related and even harder question is the location of the city center. This information is crucial for my project, but research about it on an international scale is almost non-existent. This section summarizes the choices I make when defining the set of cities for this study. Appendix A provides more details.

I start with the open collaboration database OpenStreetMap, which is based on crowd intelligence, letting users set and change geographic tags themselves. My city definition starts with all city tags contained in OpenStreetMap. I spatially join them to the Urban Centre Database

[2] `https://airdna.co`, last accessed: 16.01.2023.

[3] To the best of my knowledge, AirDNA web scraped every single property from Airbnb once every three days. This implies that a small number of properties that appeared only briefly and were immediately removed or rented might not be part of the dataset.

[4] The raw dataset contains 9,419,495 observations. However, 2,354,445 of these properties were never reserved. I have to drop another 1,917 observations because their coordinates are missing. Afterwards, I can spatially join 3,093,755 properties with my city polygons. An additional 25,603 observations drop because of missing covariates (or a missing price in 1 case). In the end, 3,068,152 entries remain.

of the Global Human Settlement Layer project (Florczyk et al., 2019).[5] In a first step, I keep all city tags within an urban center with at least 300,000 inhabitants and at least 100 Airbnbs.[6] Figure A1a shows the city delineation and the distribution of the Airbnb properties for the example of Paris.

In some cases, urban center polygons contain multiple city tags. Using the city population counts from OpenStreetMap, I retain all city tags that are associated with a population count that amounts to at least 40% of the highest population count in the respective urban center.[7] This allows me, for example, to keep Kobe, Kyoto, and Osaka as different cities without having to split the urban center around Tokyo into 120 parts. Moreover, if several tags are close to each other (less than 7 km air-line distance between the individual tags), I only retain the tag that exhibits the largest population count among them. After defining the centers (see below), I follow Akbar et al. (2021) in how I split the urban areas with more than one remaining tag, by defining border points according to their distance to the different city centers, taking the population size of these centers into account (see Appendix A). Figure A2 presents the example of Den Haag and Rotterdam, an urban center that I split into two distinct cities.

To set the location of a city tag on OpenStreetMap, users are asked to place it "at the center of the city, like the central square, a central administrative or religious building or a central road junction".[8] The coordinates of those tags therefore offer a transparent and globally consistent definition of city centers. Together with a research assistant, I manually checked all city tags that remain at this stage. Unfortunately, there are some cities for which visual inspection using satellite images and street view from GoogleMaps suggests that they are not a good representation of the actual center. I use the coordinates from OpenStreetMap for my preferred specifications whenever they seem accurate. When they do not, I use the coordinates proposed by Google Maps, which we checked as well. If they also seem not to provide a suitable representation of the center of a given city, we propose our own best guess. Figure A3 presents a corresponding example in Rosario. I provide robustness checks in which I take the coordinates from OpenStreetMap or those from Google Maps to define all city centers.

My final sample contains 734 cities worldwide. Figure 1 illustrates their geographic distribution, while Figure A4 shows the number of cities by country.

3 Implications of the monocentric city model

The basis of this analysis is the monocentric city model. It was developed by Alonso (1964), Mills (1967), and Muth (1969), who formalized an idea that had been around at least since von Thünen (1826). The model assumes an exogenously given employment center on a homogenous

[5] Their definition results in some urban areas being very broad and containing multiple well-known cities, for example Oakland/San Francisco/San José or Kobe/Kyoto/Osaka. In some of these cases setting one city center for the whole urban center would be very tricky. I therefore decided against simply adopting the definition from Florczyk et al. (2019). I believe that combining it with data on cities from OpenStreetMap results in a set of cities that is better suited for this analysis.

[6] Not including Airbnbs that have never been rented during the study period and have therefore no information about the average rent charged.

[7] I complement missing population counts on OpenStreetMap with information from Wikipedia.

[8] `https://wiki.openstreetmap.org/wiki/Tag:place%3Dcity`, last accessed: 13.01.2023.

Figure 1: 734 cities in the sample

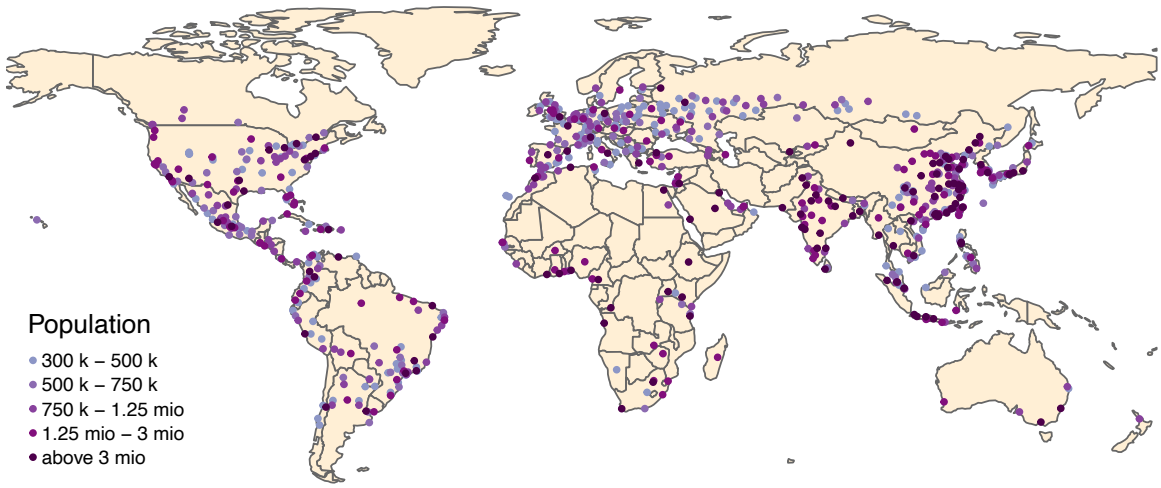

Note: The dots in this figure show the geographic distribution of the 734 cities in my sample. Their colors refer to the city's population size, with larger cities represented in darker shades. To be included in the sample, a city must have had a population of at least 300,000 inhabitants in 2015 and at least 100 Airbnb properties that were active between January 2018 and March 2019.

plain. People face transportation costs to move from their homes to the center. With very few additional assumptions, the model predicts decreasing house prices (the rent gradient), decreasing population densities (the density gradient), and increasing housing unit sizes as one moves from the city center to the periphery. The key prediction of the model is the Alonso-Muth condition:

$$R'(x) = -\frac{t'(x)}{h(x)}\ ; \tag{1}$$

where x denotes distance to the city center, $R(x)$ describes housing rents, $t(x)$ transportation costs, and $h(x)$ housing unit sizes. As the latter term describes an area, it will always be positive. While this paper deals with whether marginal transportation costs are linear, decreasing, or increasing in distance, they should certainly not be negative, which would imply that traveling longer distances is cheaper than shorter ones. Therefore both sides of eq. (1) are predicted to be negative. Economically, this signifies that housing units feature cheaper rents if they are farther away from the city center. Moreover, the rate at which housing units become cheaper in distance multiplied by their size is equal to the rate at which transportation costs become more expensive in distance. Gains through shorter transportation costs are compensated by higher housing costs, and vice versa. I use an extension of the baseline model that allows housing units to differ in size while I keep the assumption that all buildings have the same height.[9] In this

[9] Setting up the model in this way implies that flat sizes govern population density entirely. Considering a city where flat sizes increase with distance to the city center while building heights decrease, my formulation would predict even smaller flats in the central parts of the city and even larger flats towards the periphery. Allowing the height of buildings to vary is feasible from a theoretical perspective, but I lack access to appropriate data. However, due to regulations, assuming fixed building heights might not be further from reality for many cities than assuming fully flexible building heights driven by market forces.

case, the following equations hold, with $D(x)$ denoting population density:

$$D(x) = \frac{1}{h(x)} \; ; \tag{2}$$

$$R'(x) = -t'(x)D(x) \; . \tag{3}$$

The standard model predicts rents or housing prices to decrease with distance to the city center in a convex way (see Brueckner, 1987). Beyond that, the functional form of the rent gradient is not a priori clear. For (parametric) empirical estimation, however, an assumption has to be made. The most popular choices in the empirical literature are log-linear and log-log specifications, regressing log prices either on distance expressed in kilometers (or any other linear scale) or on the log of distance. The estimated parameter related to distance is called the rent gradient. The same considerations also hold for the relationship between distance and population density. For a more detailed exposition of the model see Brueckner (1987) or Duranton and Puga (2015).

Leveraging my data, I find that the log-linear assumption is more sensitive to the precise definition of cities. Figure 2 depicts this regularity for all cities with an Airbnb property further than 20 km from the city center. To construct the figure, I first recompute the rent gradients using subsets of Airbnbs within increasingly limited distances from the city center. This is equivalent to defining the cities more narrowly by imposing a maximal distance of the city fringe. I then consider the ratio between the average rent gradient for each subset and the average rent gradient computed with my preferred city definition specified below.

For the log-log specification, the cutoff distance does not alter the average rent gradient much, with a maximal deviation of 9%. However, when using a log-linear functional form, the placement of the city boundaries matters a lot. For example, suppose I limit cities to the area within 8 kilometers of the city centers. In that case, the average estimated rent gradient becomes almost 2.4 times as large compared to the case where I employ my preferred city definition. I, therefore, decide to proceed with a log-log specification.

Modeling the relations between rents and distance and between density and distance as log-log implicitly assumes that the underlying functions are

$$R(x) = Ax^b \; ; \tag{4}$$

$$R'(x) = bAx^{b-1} \; ; \tag{5}$$

$$D(x) = Cx^d \; . \tag{6}$$

For details on the derivations, see Appendix B. Plugging eq. (5) and eq. (6) into equation (3) and solving for $t'(x)$ yields

$$t'(x) = -\frac{bA}{C}x^{b-d-1} \; . \tag{7}$$

Figure 2: Functional form and city definitions

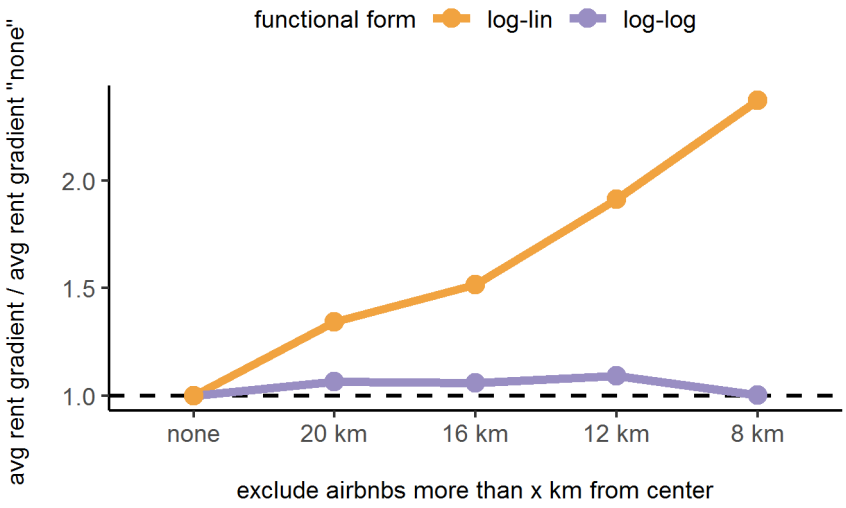

Note: This figure compares 10 average regression coefficients. In all cases, I estimate rent gradients for all cities in the sample by regressing the log of hedonic prices of Airbnbs on a distance variable interacted with city indicators. For the purple points, I use the log of the distance to the city center as the variable of interest, while I use linear distance for the orange points. I then compute the average over all city rent gradients. All regressions include city fixed effects. I apply these two ways to compute gradients for five different samples. In the baseline version, I use all Airbnb properties. This specification is denoted as "none". Within each of the two functional forms, the coefficients are divided by the average gradients of this first specification. I then apply progressively stricter thresholds, excluding all objects that are more than 20, 16, 12, and finally 8 kilometers away from the city center. This is equivalent to defining more narrow city boundaries. For this figure I only use cities that have at least one property that is more than 20 kilometers away from the city center.

Furthermore, taking the integral of this term and imposing that $t(0) = 0$ (no transportation costs to travel to the center if you are already there) returns the transportation cost function

$$t(x) = -\frac{bA}{C(b-d)}x^{b-d} \tag{8}$$

or

$$t(x) = \phi x^{\theta} \; ; \tag{9}$$

where θ describes the elasticity of transportation costs and ϕ is a constant. In other words, taking a standard version of the monocentric city model and using log-log specifications to model the rent gradient and the density gradient, as is commonly done in the empirical literature, implies that transportation costs also follow a gradient. Moreover, the transportation cost gradient θ equals the difference between the rent gradient b and the density gradient d, which, to my knowledge, is a novel result.

The rent gradient and the density gradient are both predicted to be negative by the monocentric city model; two predictions that I assess empirically below. Intuition would suggest that a

reasonable transportation cost function fulfills the following minimal requirements:[10]

1. Transportation is costly, or $t(x) > 0$ for all $x > 0$.
2. Traveling longer distances is more costly than traveling shorter distances, or $t'(x) > 0$.

As long as the predictions of the monocentric city model and therefore $b < 0$ and $d < 0$ hold, both conditions are fulfilled if $b > d$. In other words, the density gradient d needs to be steeper (more negative) than the rent gradient b. This is another prediction that I will test.[11]

Moreover, the transportation cost gradient θ holds information about the curvature of the transportation cost function. If $\theta = 1$, transportation costs are linear, or $t(x) = \tau x$. This is the typical assumption in the baseline version of the monocetric city model (Brueckner, 1987; Duranton and Puga, 2015).[12] If θ is between 0 and 1, transportation costs are concave. Several intuitive reasons would suggest this to be a plausible case. Traffic is more congested in more central locations, speed limits are lower, and buses, trains, or subways stop more frequently. Moreover, there are fixed costs associated with getting to the next public transport stop or into the car. Longer trips to locations further away (or from further away) often contain parts where traffic goes faster and the marginal costs to cover an additional kilometer decrease. If $\theta > 1$, transport costs are convex. This could, for example, be the case if more central locations are well connected but the infrastructure becomes very bad once one travels further away from the city center. In summary, I assess the following hypotheses:

- *Hypothesis 1*: Rent gradients are negative.
- *Hypothesis 2*: Density gradients are negative.
- *Hypothesis 3*: Density gradients are steeper (more negative) than rent gradients.
- *Hypothesis 4a*: Transportation costs are linear.
- *Hypothesis 4b*: Transportation costs are concave.
- *Hypothesis 4c*: Transportation costs are convex.

For each prediction, I am interested in estimating whether it holds globally and in which cities it holds to what extent.

4 Rent gradients

Before exploring the relationship between the price of a property i and its distance to the city center, it is important to account for the fact that properties in a more central location might

[10] Moreover, a person not traveling should not bear any transport costs, or $t(0) = 0$. This condition is fulfilled by construction (see Appendix B).

[11] Within this model, A and C are always positive because they are exponentials of regression intercepts (see Appendix B). When equation 8 is applied strictly, $t'(x) > 0$ is always true (as long as $b < 0$), as $b - d$ cancels. However, due to the exponentials, the constant ϕ is highly sensitive to the precise estimates. An alternative is to impose the less restrictive condition of positive transportation costs, i.e. $\phi > 0$. In that case, $t'(x) > 0$ holds if and only if $b > d$, which is the same condition as above.

[12] Linear transportation costs are a frequent assumption in all kinds of spatial models. Sometimes they are referred to as iceberg trade costs, albeit this term often refers to travels between cities.

be fundamentally different from properties further in the suburbs. Fortunately, Airbnb includes a number of variables that describe the advertised properties in an internationally standardized way.[13] In a first step, I, therefore, run the following hedonic regression of average daily rental prices on different covariates to get a measure of prices net of observable characteristics other than the distance from the city center[14]

$$\ln(\text{price})_{ic} = \alpha + \boldsymbol{\gamma}\mathbf{X}_{ic} + u_{ic} \, . \tag{10}$$

I mostly include the covariates X_i as nonparametric categories to allow for a flexible functional form. Figure A5 shows the variables I control for and their respective effects on the average daily rental price. Prices are winsorized to the 0.01 and 0.99 percentiles within each country to exclude properties with unrealistically high or low prices that risk being misreported.[15] Moreover, I normalize prices by subtracting their mean in the 734 city sample and dividing by their standard deviation. I do the same to population density below, bringing the two measures on equivalent scales. Without the normalization, the arbitrary choice of units (prices in USD or 100 USD, people per m^2 or per km^2) would influence the comparison.

As expected, the number of bedrooms and the maximum number of allowed guests both increase prices monotonically. The same is true for the number of bathrooms, exept for the highest category.[16] Hosts can charge a substantially higher price if the guests have the entire apartment for themselves, while shared rooms are cheaper than private rooms. Every additional amenity slightly increases the price, as do additional pictures of the property. Finally, prices are higher for properties close to an ocean, sea, or big lake, with a higher premium for properties situated directly at the shore.[17] The large number of observations ensures that the effects are precisely measured. I then compute the predicted log price for each Airbnb property and the residual value

$$\ln(\text{price})_{ic}^{\text{res}} = \hat{u}_{ic} = \widehat{\ln(\text{price})}_{ic} - \ln(\text{price})_{ic} \, . \tag{11}$$

As a second step, I regress these residual prices on city fixed effects and the log distance from the city center. I allow the effect of distance to vary for each city. This gives me 734 different rent gradients b_c:

$$\ln(\text{price})_{ic}^{\text{res}} = a + b_c \ln(\text{distance})_{ic} + \text{FE}_c + \varepsilon_{ic} \, . \tag{12}$$

[13] There is, of course, heterogeneity that is not described in the profile of an Airbnb property. However, I would argue that when potential renters make their decision, it is not easy for them to access information about a property unless it is provided on the platform. Nevertheless, some omitted variable bias remains, as there is information in pictures and descriptions from which I abstract in this study.

[14] One notable absence is floor area, which Airbnb does not collect. However, the number of bedrooms, the number of bathrooms, and the maximal number of guests give a good impression of the size of a property.

[15] In some cases, these high prices might also be a consequence of money laundering. Reports on money laundering on Airbnb are provided, for example, by Bell (2021) and Fazzini (2019).

[16] There are entries with an unrealistically high number of bathrooms that is probably erroneous in most cases.

[17] The indicators for proximity to the beach are the only variables that are not directly taken from Airbnb. Instead, guests can infer this distance from a map that is provided with the properties. Moreover, the hosts seem to have a clear incentive to mention a location close to a beach in the description or their photos. To construct these indicators, I measure the air-line distance from a property to the closest ocean, sea, or big lake (at least $80km^2$). To determine the location of waters, I use ESRI's "World Water Bodies" layer (`https://arcgis.com/home/item.html?id=e750071279bf450cbd510454a80f2e63`, downloaded on 10.10.2023) and the HydroLAKES data from `https://hydrosheds.org/products/hydrolakes` (downloaded on 10.01.2023).

Figure 3 shows the distribution of these gradients, both as a histogram and using a kernel density function. I estimate an average rent gradient of -0.064, implying a 0.64% decrease in price for every 10% increase in distance. The median rent gradient is -0.067, while the 25% and the 75% quartiles are -0.145 and 0.011, respectively. I find negative rent gradients for 72% of cities in my sample and statistically significantly negative rent gradients at the 5% level for 52% of cities. On the other hand, I estimate positive rent gradients for 28% of cities and statistically significantly positive rent gradients (also at the 5% level) for 13% of cities.

Figure 3: Distribution of rent gradients

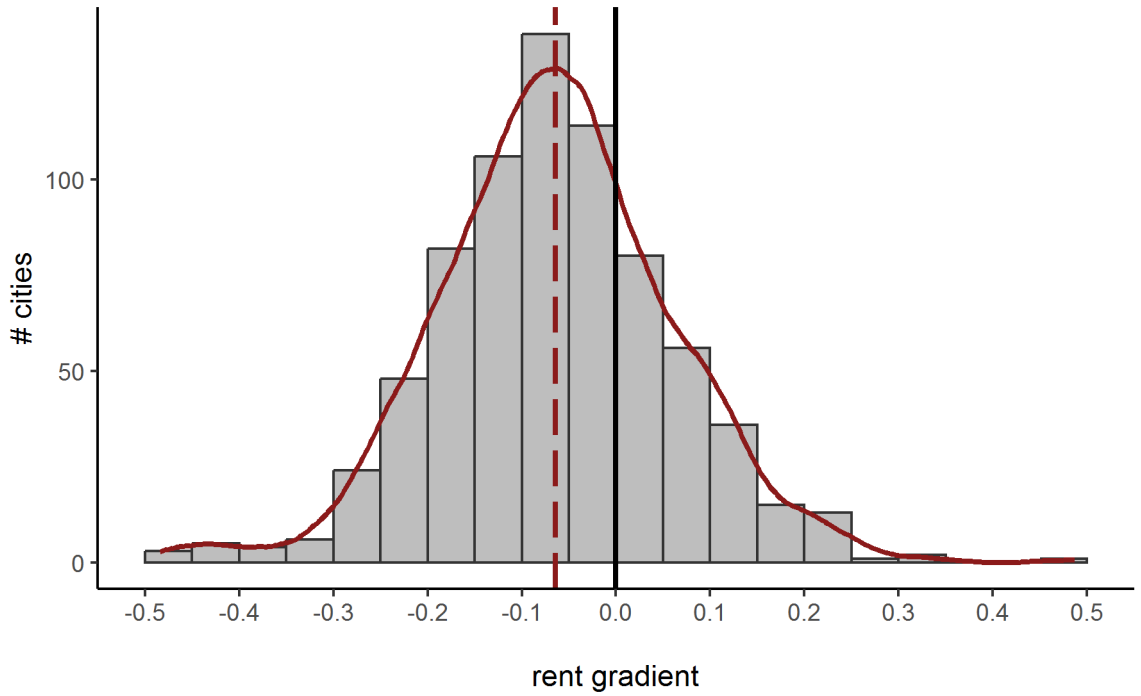

Note: This figure depicts the distribution of rent gradients for all 734 cities in my sample. I estimate these gradients using internationally consistent data from Airbnb. In particular, I regress log prices on city fixed effects, and the log distance to the city center interacted with an indicator for every city. The dashed line depicts the average rent gradient, while the solid line shows the corresponding Epanechnikov kernel density estimate. This estimation is the second stage of a two-stage procedure. The first stage is a hedonic regression in which I regress winsorized (0.01, 0.99) and standardized (subtract the mean and divide by the standard deviation) prices on different object characteristics.

Comparison with the literature

How do these estimates compare to the existing literature? Gupta et al. (2022) find a rent gradient of -0.03 and a corresponding house price gradient of -0.10 for the 30 largest MSAs in the United States.[18] The average gradient for the 70 US cities in my sample is also -0.10. Combes et al. (2018) estimate house price gradients for 277 urban areas in France. They find a median price gradient of -0.03, with the 25% quartile being -0.07 and the 75% quartile being

[18] These estimates refer to the situation before the pandemic. They find that the gradients became considerably flatter during the pandemic (0.00 / -0.09 in December 2020). This finding is confirmed by Li and Wan (2021). Their estimate of the rent gradient in Beijing flattened from -0.17 to -0.12 in June 2020 (before somewhat decreasing again). Future studies will show whether this is a pure pandemic effect that fully reverts eventually or whether this induced a more structural change.

0.01.[19] Looking at the 12 French cities included in my sample, the respective values that I find are -0.11 (median), -0.16 (25%) and -0.04 (75%). However, they find that gradients are more negative in larger urban areas, an empirical fact that I confirm below. Given that my sample only includes the very largest cities in France, this could reasonably explain my more negative estimates. Li and Wan (2021) report a pre-pandemic rent gradient of -0.17 for Beijing, with my corresponding estimate being -0.25. For Beijing, there exists also an estimate from Ding and Zhao (2014), based on housing prices between 1999 and 2003. However, they estimate a log-linear specification. Their estimates range between -0.03 and -0.07. If I apply a log-linear specification myself, I obtain an estimate of -0.03 for Beijing. Liotta et al. (2022) also apply a log-linear specification of rents on transportation costs in a robustness check.[20] They find an average estimate of -0.014 for their 192 cities, while the average log-linear estimate across the 734 cities in my sample is -0.017.

The distribution of rent gradients I find falls well within the ballpark of the existing estimates. If anything, my estimations show more negative gradients than the studies using long-term rental data while being more in line with those using house price data. This is not surprising, as long-term rents are often more tightly regulated than house prices and nightly Airbnb rates.

Comparison with long-term rental data

I can also compute rent gradients based on long-term rental properties myself and compare them with those obtained using Airbnb short-term rental properties. An advantage of this approach is that I can compare each city's gradients instead of having to rely on aggregate statistics like in the comparisons with the literature. Moreover, I can control the set of cities and define their centers and borders in a consistent manner. There are publicly available data for France and the United States, so I use these two countries for the comparison.[21] These data are based on small administrative units (block groups for the United States, communes for France) rather than on individual properties.[22]

Figure 4 shows rent gradients estimated using long-term rental data on the x-axis and corresponding gradients estimated using data on Airbnbs on the y-axis. There is a strong positive correlation between the two, although more so for France than for the US.[23] One reason for this difference could be that the French long-term rental prices are constructed from hedonic regressions, considering property-specific characteristics. The American Community Survey, on the other hand, only provides raw median rents by block group. I control for a set of covariates, but

[19] I averaged the values over their seven specifications in Table 3, Panel A.

[20] They use income net of transportation costs as the variable of interest in their main specifications, which makes it harder to compare the estimates.

[21] The French data come from *la carte des loyers*, while the US data are based on the American Community Survey. They can be found on `https://www.data.gouv.fr/fr/datasets/carte-des-loyers-indicateurs-de-loyers-dannonce-par-commune-en-2022/` (downloaded on 17.02.2023) and on `https://www.nhgis.org/` (downloaded on 16.12.2021), respectively.

[22] For the three big cities of Lyon, Marseille, and Paris, the French data contain information at the level of arrondissements.

[23] Regressing gradients obtained using Airbnb properties on gradients obtained using long-term rentals yields an coefficient of 1.12 for France, with a standard error of 0.35. For the United States, the corresponding coefficient is 0.72, with a standard error of 0.13. The underlying correlations are 0.71 (France) and 0.55 (United States), respectively.

these are also just aggregate statistics on the block group level.[24] Moreover, the US has a larger degree of local autonomy concerning taxes and public services that could influence long-term gradients beyond pure geographical considerations.

Heterogeneity

Which cities exhibit steeper rent gradients? A first factor is the size of the city. Figure 5 shows regressions of the rent gradients on the log of population size for 16 countries that are represented by more than 10 cities in my sample. I confirm the finding of Combes et al. (2018) that more populous French cities are associated with more negative gradients. Moreover, this regularity extends to other countries. It is particularly strong for European cities but also for Canada and India. While somewhat weaker, more populous cities also feature more negative gradients in the United States, China, Japan, and Malaysia. An explanation might be that location in smaller cities is less crucial, as many places are accessible within reasonable distance. This might change for larger cities, where housing location is likely to be more correlated with the part of the city in which people spend their leisure or do their shopping. Interestingly, all four countries for which bigger cities are not associated with more negative rent gradients are situated in Latin America.

I also explore heterogeneity by income level. Panel A of Figure 6 shows results for this dimension. Its base are the income groups of the World Bank.[25] I combine low-income and lower-middle-income countries into one level to get a sufficiently high number of observations.[26] For each income level, I run a separate kernel density estimation.[27] There seems to be a non-linear relation between rent gradients and income. While high-income countries have the steepest rent gradients (with an average value of -0.10), low & lower-middle-income countries exhibit similarly negative gradients (average of -0.09). Upper middle-income countries, however, have, on average, substantially flatter gradients, with an average of -0.03. The reasons for the more negative gradients will most likely differ between high income and low & middle income countries. Transportation tends to be slow in economically poorer countries (Berg et al., 2017). Car ownership is lower (Cervero, 2013) and formal public transport less expanded, while informal transport infrastructure is more common (Kumar, 2011). If the costs of commuting longer distances are higher, this might explain why people are willing to pay a higher premium to live at more central locations. In high-income countries, on the other hand, central cities might offer amenities that are nice, but not indispensable. Once people reach a certain income level, they might be willing to pay a premium to live next to these amenities.

[24] I control for the fractions of apartments that meet certain characteristics in the following categories: Number of bedrooms, number of units in the building, the year the building was built, the year the tenant moved in, presence of plumbing, presence of a kitchen and whether meals are provided, and energy source used. Moreover, I control for whether a block group is in immediate proximity to a large water body.

[25] In downloaded the income groups on 18.08.2022 from `https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups`.

[26] Low & lower middle income is still the smallest group with 135 cities, 11 of which are classified as low income by the World Bank. 336 cities are in the upper middle income category, while 262 cities are in the high income category. I drop Caracas for this graph, as the World Bank does not classify Venezuela due to issues with data availability.

[27] The area under each curve sums to one and is thus of equal size, independent of the number of observations within the group.

Figure 4: Comparisons with gradients estimated from long-term rental data

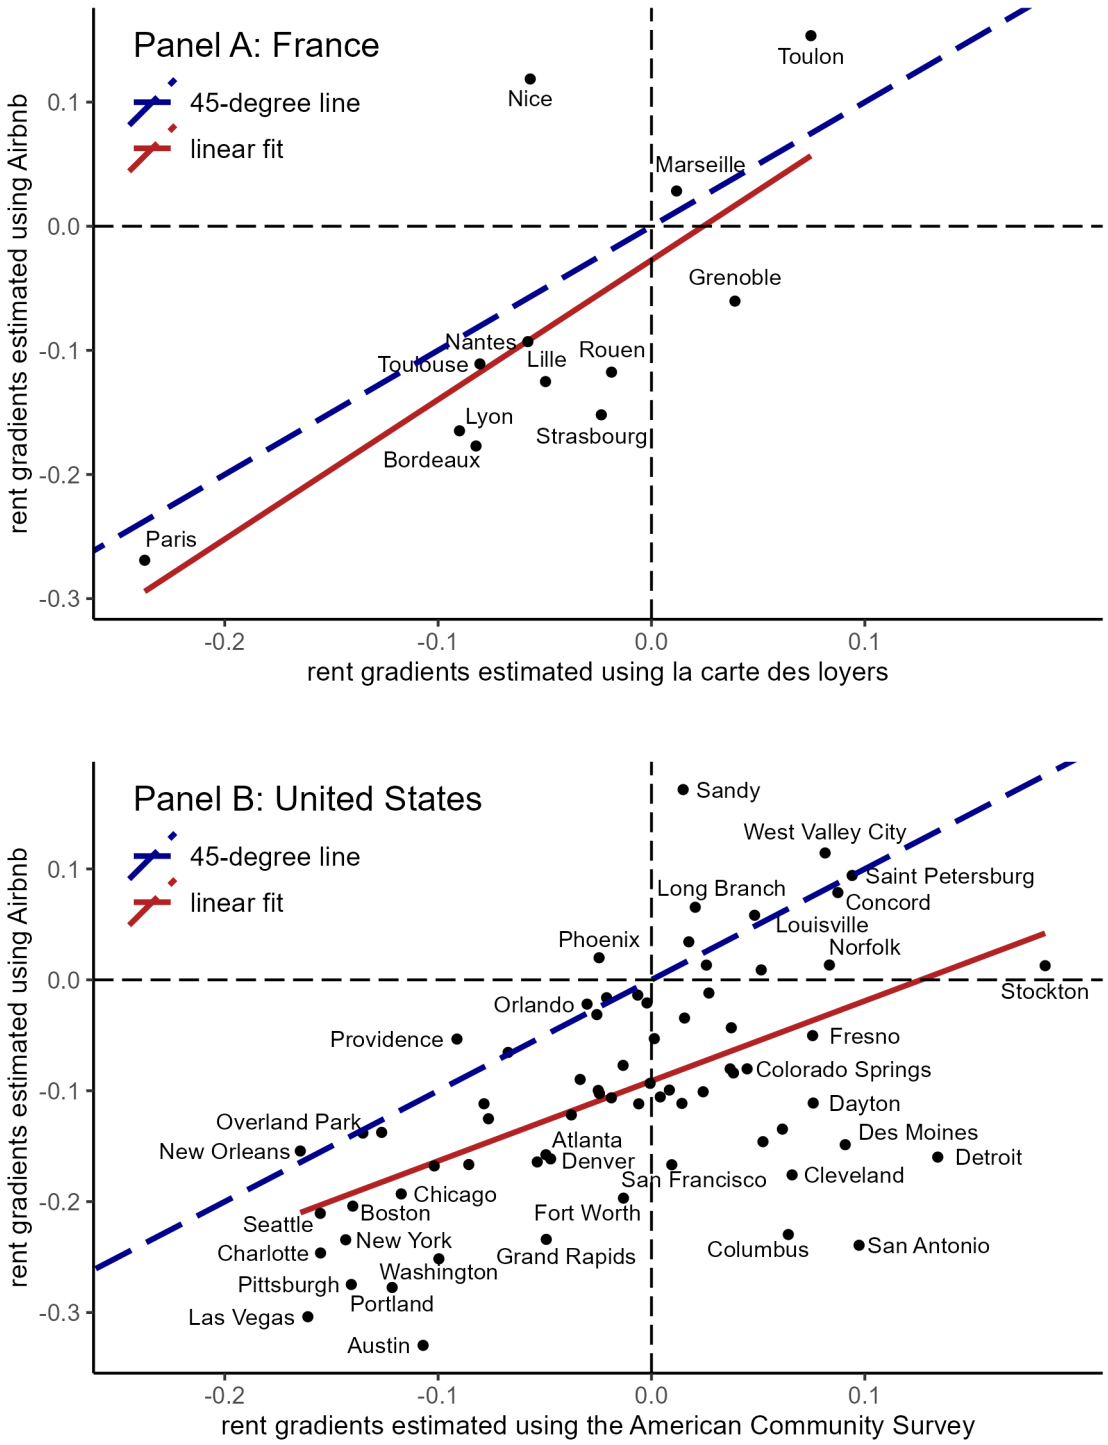

Note: Both panels show rent gradients estimated with Airbnb data on the y-axis. Panel A depicts long-term rent gradients for France on the x-axis. These gradients are estimated using hedonic prices on the commune level provided by *la carte des loyers*. The x-axis of Panel B shows long-term rent gradients for the United States. They are estimated using median rents per blockgroup from the American Community Survey (2015-2019). Unlike for France, these are raw rental prices. I therefore control for various building-characteristics at the blockgroup level. Each black point represents a city. The red line depicts the linear fit, while the blue line represents the 45-degree line.

Figure 5: Rent gradients and city size

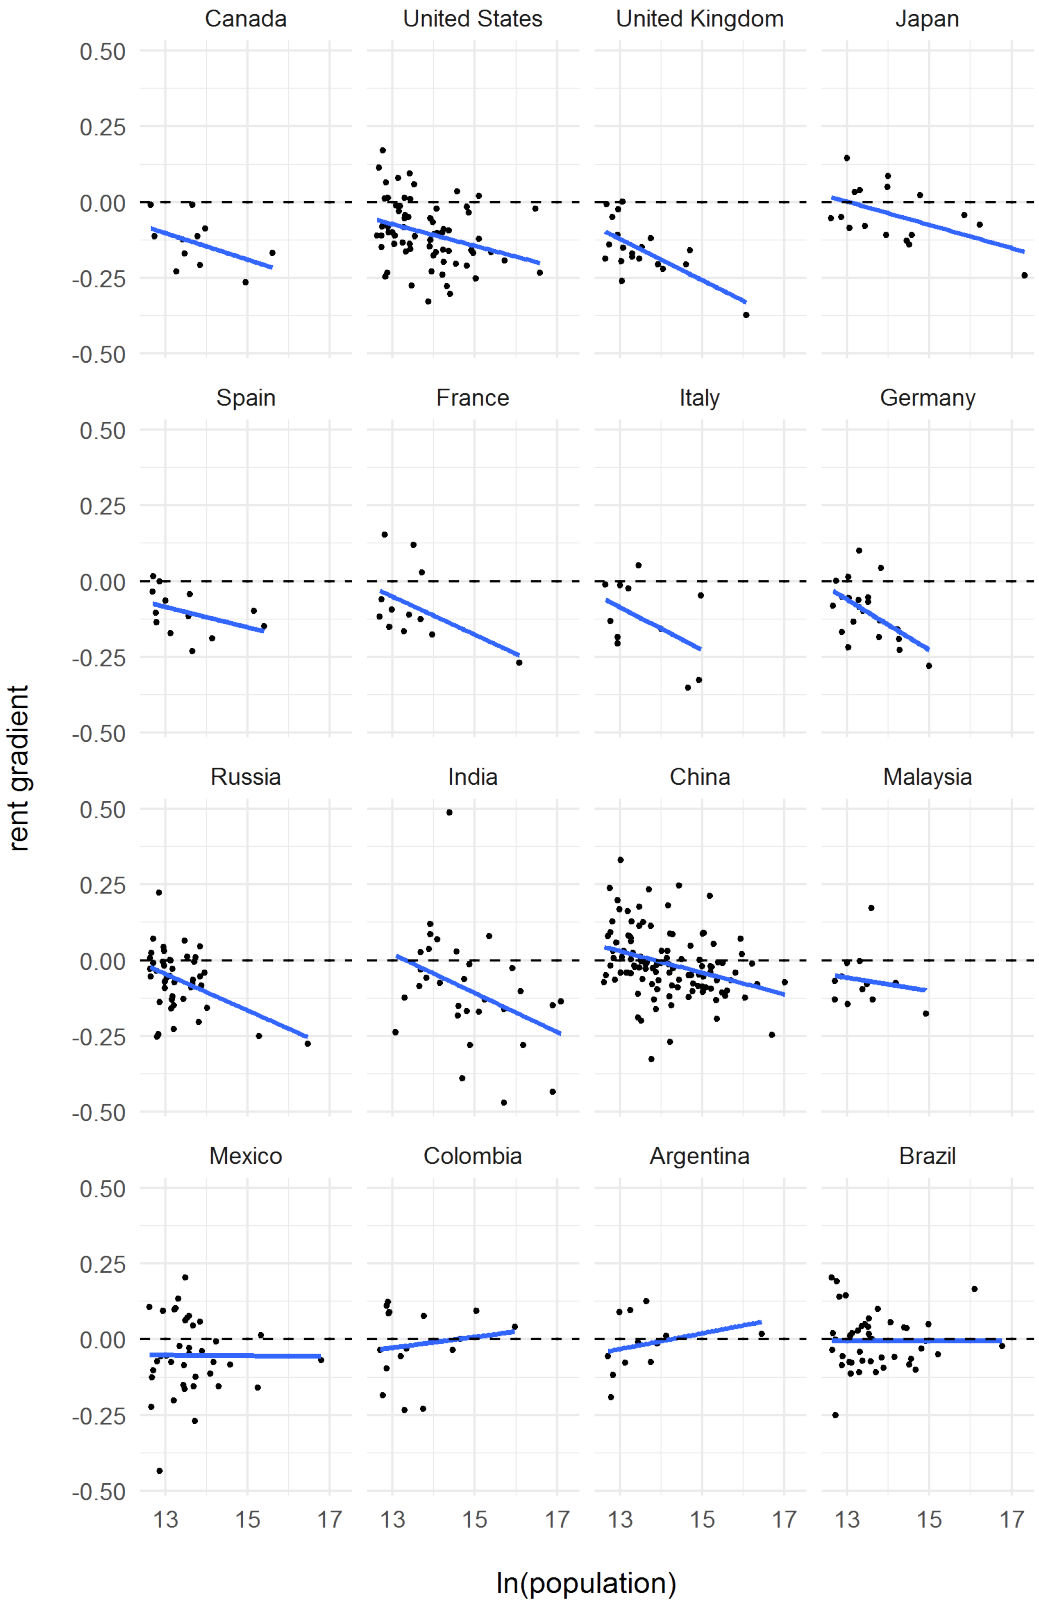

Note: The figure shows estimated rent gradients (y-axis) against the population size of cities on a log scale (x-axis). Each dot represents a city, while the blue lines show the linear fits. The comparison is shown for the 16 countries with the highest number of cities in my sample.

Figure 6: Regularities

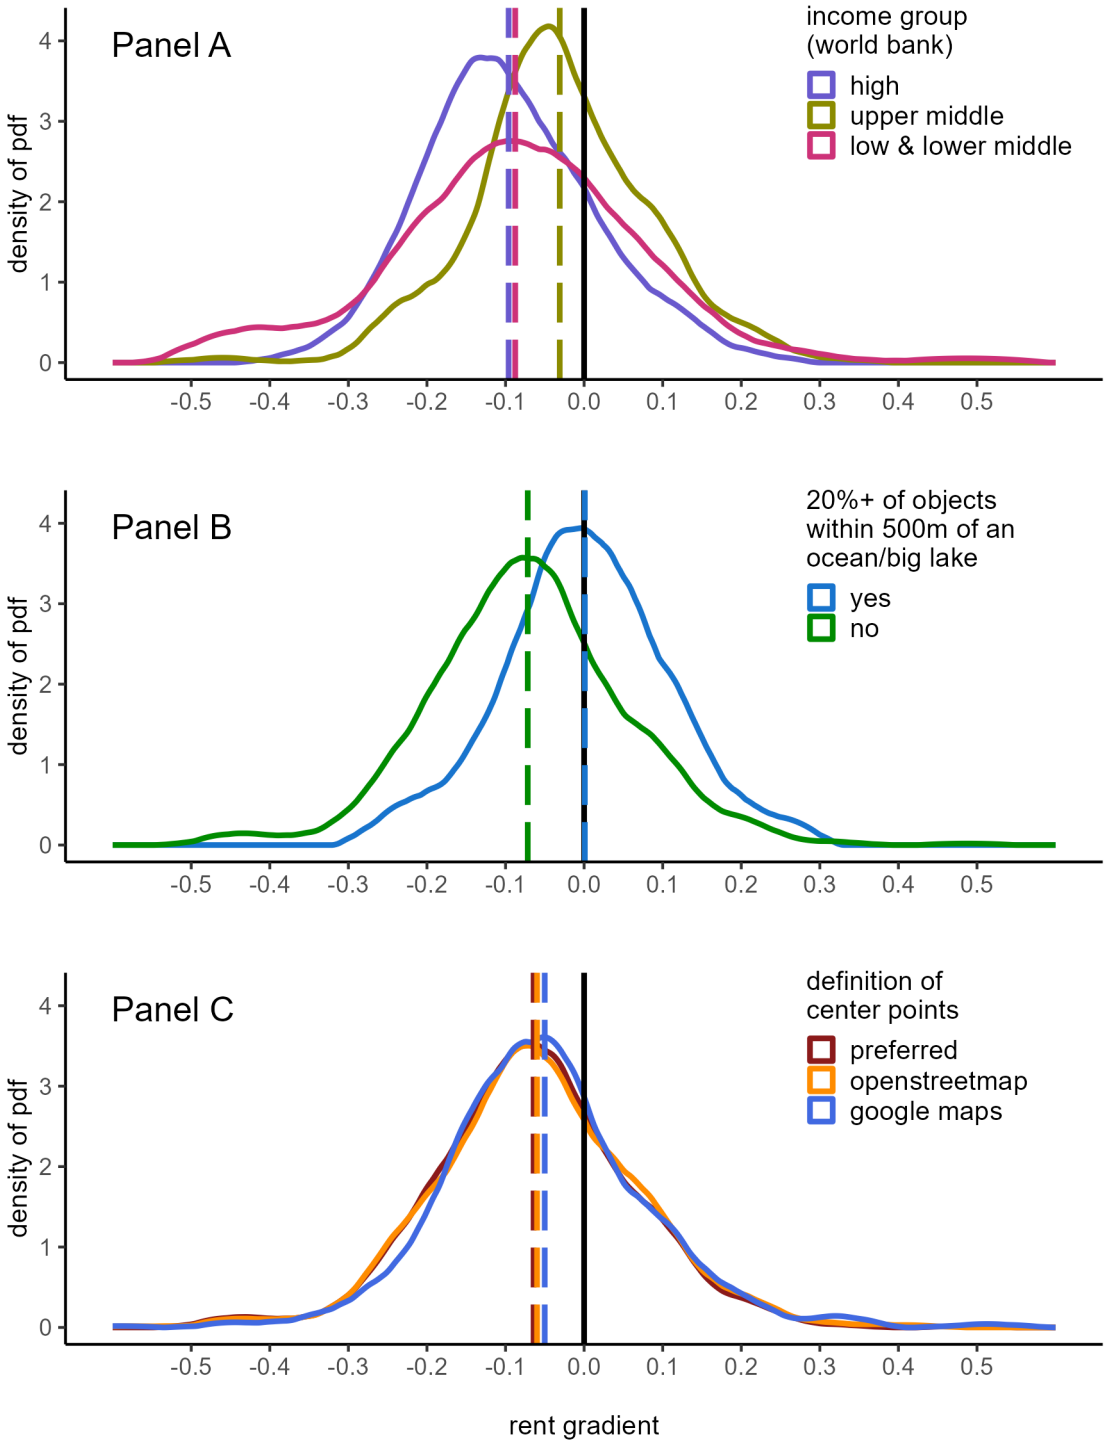

Note: Analogously to Figure 3, this figure shows distributions of rent gradients. Panel A classifies cities by their countries' income group, according to the World Bank definition. Panel B compares the subset of cities in which at least 20% of Airbnb properties are within 500m of an ocean or big lake versus those cities for which this is not the case. Panel C recomputes rent gradients using different definitions of the city center. Dashed lines represent the average rent gradients, while solid lines show the corresponding Epanechnikov kernel density estimates. All of the latter sum to one, facilitating visual comparisons between categories with different numbers of cities.

The most striking regularity concerns cities close to an ocean, sea, or big lake (at least 80 km^2). I measure this by requiring that at least 20% of Airbnbs be within 500m of the shore. I choose this relatively restrictive cutoff because I suspect that it is not the sheer presence of a large water body in proximity to a city that makes a difference, but rather whether beaches and ports play an important role in the life of its inhabitants. In a way, the seaside can act as an elongated secondary center that absorbs parts of the economic, cultural, and leisure-based activities that would otherwise be concentrated in the city center. Panel B of Figure 6 shows this analysis. The average rent gradient for seaside cities is very close to zero (the point estimate is 0.0009). Note that the first-stage hedonic regression controls for whether an individual property is in close proximity to the beach. This implies that the effect across cities is not just driven by more costly properties directly at the seaside. However, there are only 75 coastal cities in the sample, according to the definition above. Moreover, if one expects Airbnb prices to be a bad proxy for other types of real estate prices, despite the evidence presented above, this would probably be the most vulnerable sub-analysis, as access to beaches might be valued more highly by tourists than by permanent residents, and Airbnb properties might be clustered particularly close to the seaside. Nevertheless, the difference is sizeable and supported by the fact that Liotta et al. (2022) also report a weaker relationship between rents and transportation costs for coastal cities.[28]

Robustness check

Panel C of Figure 6 revisits the question of the optimal placement of the city center. It shows three different kernel density functions that differ in the source of the center coordinates. The red curve represents my preferred center definitions described in Section 2. The orange curve takes the city coordinates from OpenStreetMap for all cities, while the blue curve does the same with the city tags from Google Maps.[29] Choosing a different center source does not fundamentally alter the rent gradient distribution. In particular, taking all centers from OpenStreetMap results in a similar distribution as my preferred center choices. The respective average values are -0.064, -0.060, and -0.050. The fact that my preferred specification results in the most negative gradients is consistent with it having the least measurement error and therefore being the least biased towards zero.[30]

[28] In their case, they do not consider transportation costs per se, but income net of transportation costs.

[29] To pin down the center locations proposed by Google Maps, I searched for the route to travel to a given city and chose the endpoint. These are also the places on top of which Google Maps depict the city names.

[30] In that regard, the fact that using Google Maps returns an average estimate that is closest to zero is also in line with my expectations. Together with a research assistant, I evaluated the center candidates from OpenStreetMap and Google Maps on a scale from 1 (suitable choice for the center) to 3 (absolutely not suitable choice for the center). Whenever the center from OpenStreetMap is classified as a 1, it is locked as my preferred center choice. If it is classified as a 2 or a 3 and the center inferred from Google Maps is classified as a 1, I take the latter as my preferred center. We only manually defined an alternative when both sources were classified as unsuitable for the center. Center candidates from OpenStreetMap are classified as 1 in 76% and 3 in only 4% of cases. In comparison, the candidates from Google Maps are classified as 1 in 59% of cities and 3 in 14%. According to our assessment, the centers inferred from Google Maps are, therefore, less accurate (at least, this was the case at the time of the analysis). This is consistent with a larger bias towards zero.

5 Density gradients

Data on population density are taken from the Global Human Settlement Project described in Section 2 and Appendix A. I use the GHS-POP file (Schiavina et al., 2019) that is computed by combining population data from administrative sources and machine-learning-based detection of artificial structures. I employ the version with the 250m × 250m resolution.[31] Figures A1b (Paris), A2b (Den Haag and Rotterdam), and A3b (Rosario) show population counts on this resolution for exemplary cities. I drop grid cells with a value of one inhabitant or less. These are mostly located in water bodies or other areas not suited or developed for housing.[32] Moreover, I again normalize population density by subtracting its mean and dividing by its standard deviation to bring it on the same scales as rents and make the estimated gradients as comparable as possible. I then estimate

$$\ln(\text{population density})_{ic} = c + d_c \ln(\text{distance})_{ic} + \text{FE}_c + \varepsilon_{ic} \,, \tag{13}$$

where FE_c are city fixed effects, and d_c denotes one density gradient per city.

Figure 7 shows the distribution of density gradients in the sample. I estimate an average density gradient of -0.356. If the distance to the city center increases by 10%, density decreases by 3.56% on average. The median, 25% quartile, and 75% quartile are -0.344, -0.494, and -0.203, respectively. Concerning hypothesis 2, I find density gradients to be negative in 96% of the cities in my sample and statistically significantly negative at the 5% level in 93% of cities.

There exist positive gradients as well. This is true for 4.1% and statistically significantly so for 3.5% of the cities in my sample (also at the 5% level). I can confirm the finding in the literature that South African cities have positive gradients, explained by discriminatory housing rules during apartheid (Selod and Zenou, 2001). I estimate positive gradients for all six South African cities included in this study.

6 Transportation cost gradients

As derived in Section 3, under certain functional-form assumptions, the transportation cost gradient θ is equal to the difference between the rent gradient b and the density gradient d. In order to obtain transportation costs that are positive and increasing with distance, d needs to be more negative (steeper if $d < 0$ and $b < 0$) than b. Having estimated rent gradients and density gradients for all 734 cities, I can now assess hypothesis 3. I obtain a positive $\hat{\theta}$ in 89% of the cases. Only for 11% of the cities is the density gradient larger (less negative) than the rent gradient. This includes cities for which I have estimated a positive $\hat{b}$ and/or $\hat{d}$, which makes them incompatible with the theoretical predictions of the monocentric city model in the first place. The subset with negative values for both $\hat{b}$ and $\hat{d}$ contains 513 cities. However, within this subset, I also obtain a positive $\hat{\theta}$ for 89% of cities.

[31] For the delineation of cities I keep the 1km × 1km grid size on which the GHS urban center database is constructed.

[32] To see why this is necessary, consider the case of New York City. Manhatten is, without a doubt, very densely populated. It also hosts the city center. However, it is surrounded by water. Including the pixels located in the water would lead to a severe underestimation of the density of the areas around the city center.

Figure 7: Distribution of density gradients

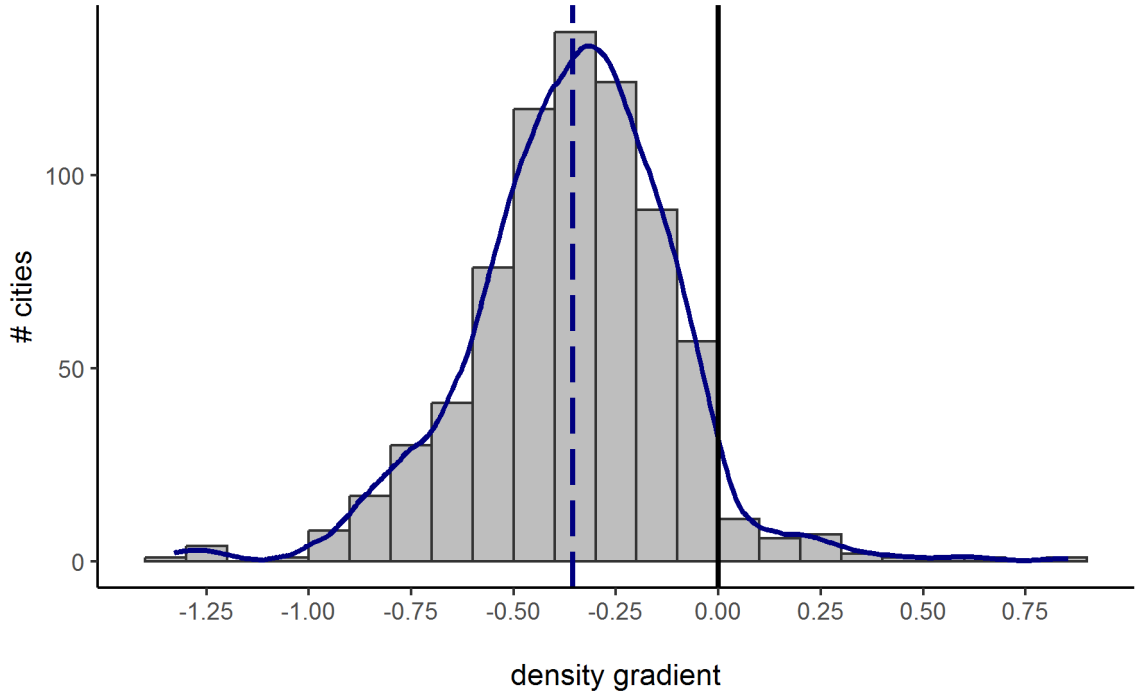

Note: This figure shows the distribution of density gradients for all 734 cities in my sample. In particular, I regress log population counts on city fixed effects, and the log distance to the city center interacted with an indicator for every city (shown here). The dashed line depicts the average density gradient, while the solid line shows the corresponding Epanechnikov kernel density estimate. The underlying data come from the Global Human Settlement (GHS) project. They disaggregate administrative population counts to 250m × 250m cells using the fraction of built-up area in a cell. I standardize the data by subtracting their sample mean and dividing it by their standard deviation to facilitate the comparison with the rent gradients.

Figure 8 shows the distribution of $\hat{\theta}$. Figure A6 provides an equivalent representation for the subset of cities for which I estimate negative rent and density gradients. The results are qualitatively similar. I estimate a mean transportation cost gradient of 0.29. The median, the 25%, and the 75% quartile are 0.28, 0.13, and 0.44, respectively.

The dotted line represents a transportation cost gradient of 1. This is the usual assumption in baseline versions of the monocentric city model model. However, based on my estimation, I reject linear transportation costs for the overwhelming majority of cities. Instead, it suggests that we should think about transportation costs as a concave function in distance.

While I estimate the transportation cost function to be concave for most cities, there is substantial heterogeneity concerning the curvature. In other words, how fast marginal transportation costs decrease in distance is very different across cities. Table 1 depicts the average $\hat{\theta}$ by country or region.[33] There appear to be some general trends: the US, Canada, and the United Kingdom are all on the lower end of the spectrum, while Latin American countries, as well as France, Italy, and Spain exhibit steeper transportation cost gradients.

[33] All countries with at least 10 cities in the sample are included as such. All other countries are assigned to geographical regions following the definition of the World Bank, downloaded on 2022-08-18 from `https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups`.

Figure 8: Distribution of transportation cost gradients

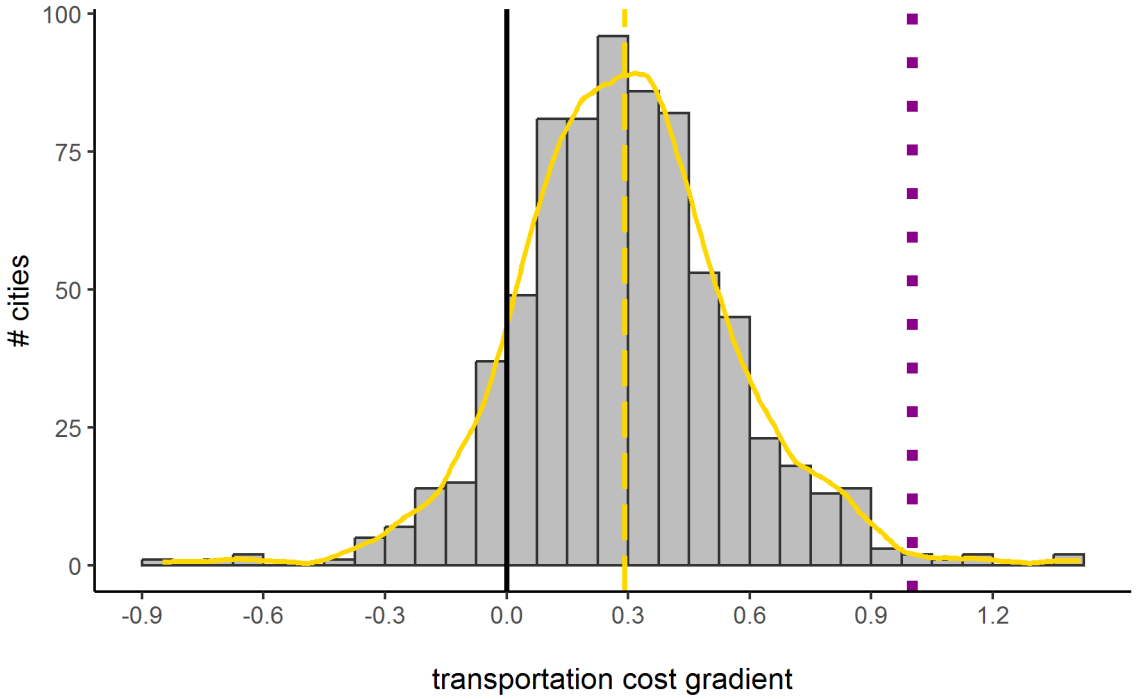

Note: This figure depicts the distribution of transportation cost gradients for my 734 sample cities. As theoretically suggested by the monocentric city model with log-log relations, these gradients are inferred from the difference between rent and density gradients. The dashed line depicts the average inferred transportation cost gradient, while the solid line shows the corresponding Epanechnikov kernel density estimate. The dotted line indicates an elasticity of transportation costs of one, equivalent to linear transportation costs.

Figure A7 translates these elasticities of transportation costs with regards to distance to transportation cost functions. It includes all countries with more than 10 cities in the sample. For each of these countries, I compute the average of the estimated $\hat{\theta}$ and use it to calibrate the transportation cost function $t(x) = \phi x^{\theta}$. I set ϕ to one for all countries. Hence, the figure does not imply that France has higher inferred transportation costs than the US for any given distance. Instead, the curves can be interpreted as transportation costs by distance relative to the transportation costs for a one kilometer trip. Thus, the figure suggests that the transportation costs of the first kilometer of transportation account for a larger part of transportation costs in the United States than in France. Possible explanations for this include higher fixed costs from getting to the bus stop and waiting for the bus, or a high level of congestion at central locations, combined with free-flowing traffic further away from the center.

7 Discussion

By combining the rent gradient b and the density gradient d, I provide an estimate for the transportation cost gradient θ. If I write equation 9 in logs, I obtain

$$ln(t(x)) = \Phi + \theta ln(x) \ , \qquad \text{where } \Phi = ln(\phi) \ , \tag{14}$$

Table 1: Average inferred transportation cost gradients by country / region

	average $\hat{\theta}$	# cities
United States	0.07	70
Sub-Saharan Africa	0.12	37
Canada	0.13	11
United Kingdom	0.14	20
Malaysia	0.16	12
Germany	0.17	21
Russia	0.24	44
Mexico	0.25	38
Brazil	0.25	44
Europe & Central Asia	0.31	74
India	0.33	31
Japan	0.33	17
Latin America & Caribbean	0.35	40
East Asia & Pacific	0.35	50
Middle East & North Africa	0.35	30
South Asia	0.36	7
Argentina	0.36	12
Spain	0.36	13
Colombia	0.40	16
Italy	0.41	11
Indonesia	0.45	11
China	0.45	113
France	0.47	12

with θ being the elasticity of transportation costs to living further away from the city center. This equation can, in principle, be directly estimated, even if it is not easy to find good data on transportation costs, in particular, if the choice of the transport mode is relevant.

Crucially, the $\hat{\theta}$ estimated in this paper is closely related to the parameter of urban costs that is an input in the model of Duranton and Puga (2022).[34] In their model, they interpret the parameter as the "elasticity of urban costs with respect to city population" (p. 3). They provide three distinct empirical ways to estimate this elasticity. One of these ways is closely related to the methodology deployed in this paper. The most notable difference is that I allow for different housing unit sizes at different locations within a city, which leads me to include data on density in the estimation. Their work focuses on the United States, and all their estimation strategies lead them to a parameter of 0.07. My average $\hat{\theta}$ for all United States cities included in my sample precisely matches this value.[35] This is reassuring concerning the validity of the Airbnb data for questions beyond the short-term rental market itself.

However, according to my analysis, the United States' $\hat{\theta}$ is at the low end of the spectrum. The global average urban cost estimator might be closer to 0.29 or about four times as large. Future research is needed to assess whether the other empirical strategies used by Duranton and Puga

[34] The corresponding parameter in their model is denoted γ.

[35] They provide just one estimate from a pooled regression of the entire country. Given my estimates, I consider it plausible that their data and methodology would also lead them to negative values for certain cities if they would allow their parameter to vary by city.

(2022) also show a similar pattern when applied to other countries. Furthermore, even within the US, their pooled estimate masks a distribution of patterns, and urban costs might vary considerably for different cities.

Figure 9 depicts this heterogeneity among cities within the United States.[36] It also suggests that the prevalent transportation modes in a city matter for the elasticity of transportation costs. Using data from the American Community Survey, I compute the share of commutes done by foot, bike, or public transport. Without claiming causality, this share appears to be correlated with $\hat{\theta}$. This relation makes intuitive sense. When looking at trips by bike or walking, a linear transportation cost function suddenly seems not too unrealistic anymore, as we are limited by our physical capacities. Moreover, public transport networks are often quite dense and well-served in more central parts of cities, while it becomes more cumbersome to travel by public transport in the suburbs, as stops are further apart and connections are less frequent. On the other hand, these are places where traveling by car is very convenient due to lower congestion levels.

Figure 9: US heterogeneity and transport modes

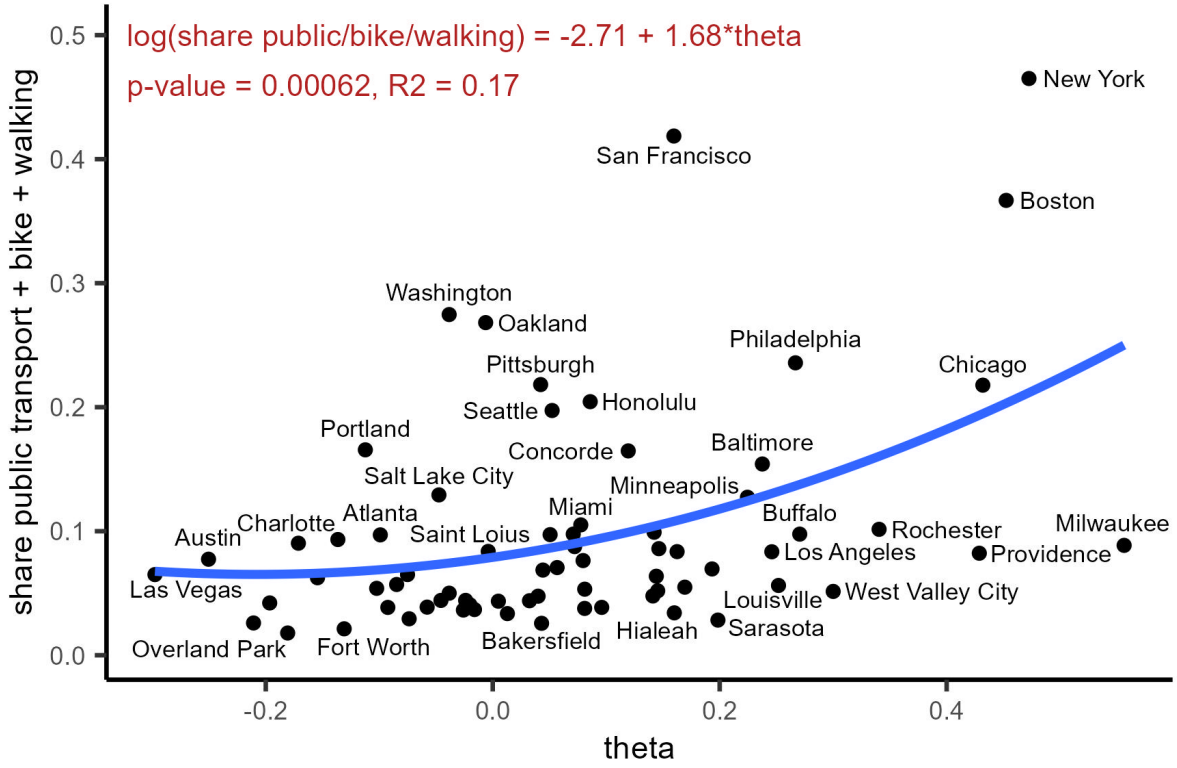

Note: For the 70 US cities in my sample, this figure shows the relation between the transportation cost gradient and the share of people commuting by public transport, bike, or walking. The x-axis shows the inferred transportation cost gradients $\hat{\theta}$ that are computed from the difference between rent and density gradients. The y-axis shows the share of people using public transport, bikes, or walking for their commute to work. The corresponding data come from the American Community Survey (2015-2019). When computing that share, people working from home were disregarded from the calculation of the total, as were answers referring to unspecified "other means".

[36] The US has particularly many negative $\hat{\theta}$, consistent with the low average $\hat{\theta}$. However, it also exhibits the lowest $\hat{\theta}$ when focusing on the subset of cities with negative rent and density gradients.

Abstracting from the formal model of Duranton and Puga (2022), a more literal interpretation of θ also holds interesting insights. An exponent of 0.07 suggests that marginal transportation costs decrease extremely quickly with distance. Living far beyond the most densely populated areas comes with very limited additional private transport costs. An exponent of 0.30 implies that living in more remote locations increases private transport costs by much more. This can have implications for urban planning. Imagine there is a consensus to densify cities, in order to counteract negative social externalities from a more dispersed population. If the private costs of travelling additional distance are lower, measures will have to be stronger and affect more people, in order to achieve the same effect.

8 Conclusion

I use novel data on over 3 million short-term rental properties from Airbnb to estimate rent gradients in a sample of 734 cities worldwide. The resulting estimates are similar to those of the few existing studies computing gradients using house prices or long-term rents in a few select countries. I find rent gradients to be negative for most cities, with an average elasticity of -0.064 between price and distance. However, there are a substantial number of cities for which estimated rent gradients are flat or even positive. These are, in particular, cities that are smaller, situated at the shore of a large water body, or located in upper-middle-income countries. I also estimate density gradients for the same set of cities. Density gradients are steeper/more negative than rent gradients in almost 90% of cities, with an average elasticity of -0.36.

My results suggest that computing these gradients using a log-log specification is less sensitive to the precise city definitions than the alternative log-linear specification that is also much used in the literature. I show that imposing log-log gradients to a standard version of the monocentric city model implies a third gradient of transportation costs as a function of the distance to the city center. This gradient equals the difference between the rent and the density gradient. I estimate it to be 0.29 on average, which implies concave transportation costs. For almost all cities, I can reject the assumption of linear transportation costs (elasticity of 1).

The transportation cost gradient maps to the urban cost parameter in the model of Duranton and Puga (2022). They focus on the United States, and I closely match their estimate using the US cities in my sample. However, my research suggests that the United States is an outlier on the global scale, with the worldwide average transportation cost gradient being about four times as large as that of the United States. Moreover, there is also considerable heterogeneity among cities within the US. In particular, cities with a higher share of commutes by foot, bike, or public transport tend to have transportation costs that are more heavily influenced by distance.

A promising path for future work will be to directly estimate the transportation cost gradient empricially and to compare it with the transportation cost gradient that can be inferred from the monocentric city model and the rent and density gradients. The work by Akbar et al. (2021), who measure car travel transportation costs in urban India using Google Maps, could be an exciting starting point. The increase in data availability might also make it possible to take other modes of transportation into account. It would be interesting to observe whether directly

estimated transportation cost gradients can be squared with the assumptions and predictions of the monocentric city model.

While highly stylized, the monocentric city model remains relevant to this day. It helps us to understand the interdependence of transportation costs, population density, and real estate prices. This paper shows that a number of its key predictions still hold for most large cities worldwide.

9 Bibliography

Akbar, P. A., Couture, V., Duranton, G. and Storeygard, A. (2021). Mobility and congestion in urban India. Working Paper.

Almagro, M. and Domínguez-Iino, T. (2022). Location sorting and endogenous amenities: Evidence from Amsterdam. Working Paper.

Alonso, W. (1964). *Location and Land Use. Toward a General Theory of Land Rent*. Harvard University Press.

Barron, K., Kung, E. and Proserpio, D. (2021). The effect of home-sharing on house prices and rents: Evidence from Airbnb. *Marketing Science* 40: 23–47.

Bell, A. (2021). Money Laundering May Be Easier in the Digital World. In Idzikowski, L. (ed.), *Money Laundering*. Greenhaven Publishing, 67–71.

Berg, C. N., Deichmann, U., Liu, Y. and Selod, H. (2017). Transport policies and development. *The Journal of Development Studies* 53: 465–480.

Brueckner, J. K. (1987). The structure of urban equilibria: A unified treatment of the Muth-Mills model. In Mills, E. S. (ed.), *Handbook of Regional and Urban Economics 2*. Elsevier, 821–845.

Calder-Wang, S. (2021). The distributional impact of the sharing economy on the housing market. Working Paper.

Cervero, R. (2013). Linking urban transport and land use in developing countries. *Journal of Transport and Land Use* 6: 7–24.

Coles, P., Egesdal, M., Ellen, I., Li, X. and Sundararajan, A. (2018). Airbnb usage across New York City neighborhoods: Geographic patterns and regulatory implications. In Davidson, N. M., Finck, M. and Infranca, J. J. (eds), *The Cambridge Handbook of the Law of the Sharing Economy*. Cambridge University Press, 108–128.

Combes, P.-P., Duranton, G. and Gobillon, L. (2018). The costs of agglomeration: House and land prices in French cities. *The Review of Economic Studies* 86: 1556–1589.

Ding, C. and Zhao, X. (2014). Land market, land development and urban spatial structure in Beijing. *Land Use Policy* 40: 83–90.

Duranton, G. and Puga, D. (2015). Urban land use. In Duranton, G., Henderson, J. V. and Strange, W. C. (eds), *Handbook of Regional and Urban Economics 5*. Elsevier, 467–560.

Duranton, G. and Puga, D. (2022). Urban growth and its aggregate implications. Working Paper, CEPR.

Fazzini, K. (2019). How criminals use Uber and Airbnb to launder money stolen from your credit card. *CNBC* Available at: `https://www.cnbc.com/2019/02/07/how-criminals-use-airbnb-uber-launder-stolen-credit-card-money.html` (last accessed: 15.06.2023).

Florczyk, A., Corbane, C., Schiavina, M., Pesaresi, M., Maffenini, L., Melchiorri, M., Politis, P., Sabo, F., Freire, S., Ehrlich, D., Kemper, T., Tommasi, P., Airaghi, D. and Zanchetta, L. (2019). GHS urban centre database 2015, multitemporal and multidimensional attributes, R2019A. European Commission, Joint Research Centre (JRC). Dataset.

Fuentes, E. (2019). Why GPS coordinates look wrong on maps of China. `https://www.serviceobjects.com/blog/why-gps-coordinates-look-wrong-on-maps-of-china/` (last accessed: 13.01.2023).

Garcia-López, M. Àngel, Jofre-Monseny, J., Martínez-Mazza, R. and Segú, M. (2020). Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona. *Journal of Urban Economics* 119: 103278.

Gupta, A., Mittal, V., Peeters, J. and Van Nieuwerburgh, S. (2022). Flattening the curve: Pandemic-induced revaluation of urban real estate. *Journal of Financial Economics* 146: 594–636.

Kumar, A. (2011). Understanding the emerging role of motorcycles in African cities : A political economy perspective. Sub-Saharan Africa transport policy program (SSATP) Discussion Paper, World Bank Group.

Li, L. and Wan, L. (2021). Understanding the spatial impact of COVID-19: New insights from Beijing after one year into post-lockdown recovery. Working Paper, SSRN.

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

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

Muth, R. F. (1969). *Cities and Housing*. University of Chicago Press.

Schiavina, M., Freire, S. and MacManus, K. (2019). GHS population grid multitemporal (1975, 1990, 2000, 2015) R2019A. European Commission, Joint Research Centre (JRC). Dataset.

Selod, H. and Zenou, Y. (2001). Location and education in South African cities under and after apartheid. *Journal of Urban Economics* 49: 168–198.

von Thünen, J. H. (1826). *Der Isolierte Staat in Beziehung auf Landwirtschaft und Nationalökonomie*. Perthes.

A City definitions and delineations

This appendix describes the choices I make when defining the cities that form the basis of the analysis in a more detailed way.

1. I start with all city tags from OpenStreetMap.[37] I downloaded all entries with a place = city tag using Overpass turbo.[38] At the time of the download, OpenStreetMap contained 10,394 city tags worldwide.

2. The Global Human Settlement Layer project of the European Commission provides the Urban Centre Database UCDB R2019A.[39] Their definition of urban areas mainly builds on two factors: i) built-up area, evaluated from (daylight) satellite data using machine learning techniques and ii) administrative population data. 1km × 1km grid cells with an estimated population of at least 1,500 or a built-up area of at least 50 % form the basis of their 13,135 urban areas. According to their estimation, 1,799 of the urban areas were inhabited by at least 300,000 people in 2015. Additionally, I filter out urban areas for which I do not have data on at least 100 Airbnbs, including data on prices charged for at least one night over the study period. This further lowers the number of urban areas to 721.

3. I spatially join the city tags to the urban areas. After this step, there are 2010 tags within 707 urban areas remaining. For these tags, I look at the population count that is linked to them in OpenStreetMap. Whenever this information is not available (451 cases), I try to add it from Wikipedia,[40] using the English version whenever possible, but resorting to other languages if the English version has no population count. In 11 cases this still does not yield a result. However, all of these tags either represent subcenters that appear to be a lot smaller than another city in the same urban area, or they are in close proximity to another tag that represents essentially the same city. I identify the highest population count among all city tags in an urban area and remove the tags that have a population count of less than 40% from that maximum. After this step, 852 cities remain.

4. For all of these cities, the accuracy of the center coordinates proposed by 1) OpenStreetMap and 2) Google Maps[41] was visually assessed. For 76% of cities the coordinates from OpenStreetMap provide a very accurate location. In another 8% of cities I resort to the coordinates proposed by Google Maps instead. For the remaining 16% of cities, I provide an own best guess. For Chinese cities, the maps displayed by Google Maps are not superimposable to satellite images because of government regulations. Instead they are shifted in a non-monotonic way (Fuentes, 2019). Airbnb uses Google Maps to display the location of their properties. As I am eventually interested in the distance of Airbnb properties from the city center, this is consistent with the center coordinates from Google Maps. However, the coordinates provided by OpenStreetMap appear to be based on the

[37] `https://openstreetmap.org` I last accessed this and all other websites in this section on 12.01.2023.
[38] `https://overpass-turbo.eu`, downloaded on 25.06.2021.
[39] `https://ghsl.jrc.ec.europa.eu/ghs_stat_ucdb2015mt_r2019a.php`, downloaded on 30.01.2019.
[40] `https://wikipedia.org`.
[41] `https://google.com/maps`.

actual satellite images. I therefore artificially "falsify" their center coordinates by shifting them in a way that makes them consistent with Google Maps.

5. Once the city centers are determined, I split urban centers with multiple city tags. However, there are cases in which multiple tags are so close that it is unlikely that they constitute two separate cites, in which case I keep them as one city. To filter out these cases, I measure the distances between all tags in an urban area. If they are less than 7 kilometers apart I classify them as neighbors. I then use network analysis to find components, that is, sets of neighbors that are directly or indirectly (a neighbor's neighbor) connected to each other, but not to any other city tag. Within each component I only keep the tag with the largest population. The number of urban areas is still unchanged at 707 after this step and so is their extent. However, the number of cities deacreases to 800.

6. In cases in which an urban area is located in one single country, I use the rule described by Akbar et al. (2021) (Appendix A, point 7) to split urban areas that host multiple cities. Keeping the 1km × 1km grid structure of the GHSL, I compute the distances of each grid cell centroid to the different city centers. To split two cities A and B, border points X are assigned such that

$$\frac{\text{dist}(X,A)}{\text{dist}(X,B)} = \left(\frac{\text{Pop A}}{\text{Pop B}}\right)^{\frac{0.57}{2}} \tag{15}$$

   where $\text{dist}(X,A)$ denotes the distance of a grid point X to the center of city A and dist (X,B) denotes the distance of the grid point the the center of city B. In some cases, this creates little enclaves; city parts that are not connected to the rest of the city. If the enclaves have only one city they share a border with (defined as sharing at least one edge of one grid cell) they are reassigned to that city. This already solves most cases. The procedure is then repeated until all enclaves are reassigned. Urban areas that span across a national border are split at the border. In this case, enclaves that do not contain a city tag are disregarded.

7. In a final step, I recount the number of Airbnbs in each newly defined city. I also recompute the number of inhabitants, based on the GHS-POP file from the Global Human Settlement Layer project (Schiavina et al., 2019).[42] Consistent with the rule used above, I discard cities with less than 300,000 inhabitants or for which I have information on less than 100 Airbnb properties. My final sample contains 734 cities worldwide.

[42] `https://ghsl.jrc.ec.europa.eu/download.php?ds=pop`, downloaded on 06.08.2021.

Figure A1: Example of city definition: Paris

(a) Distribution of Airbnb properties

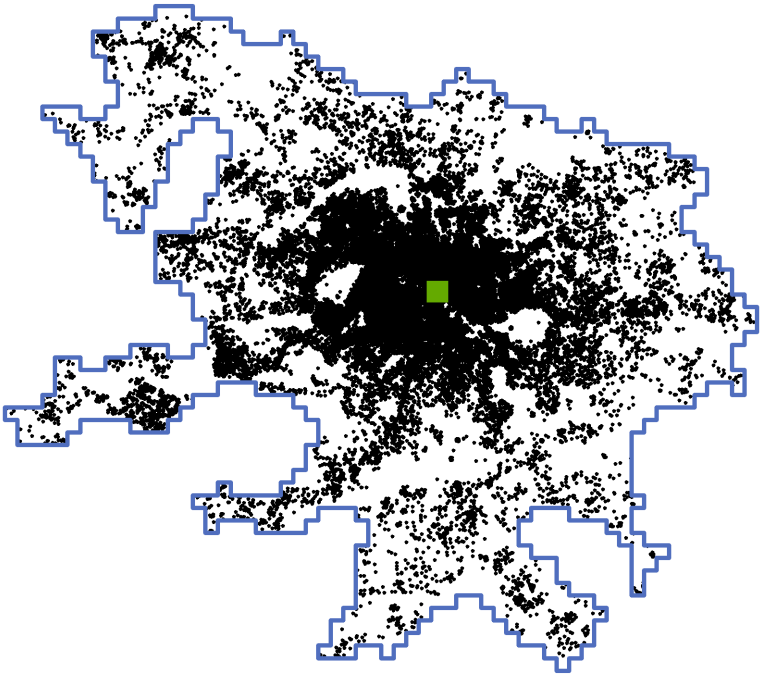

(b) Population count in 250 × 250 meter grid cells

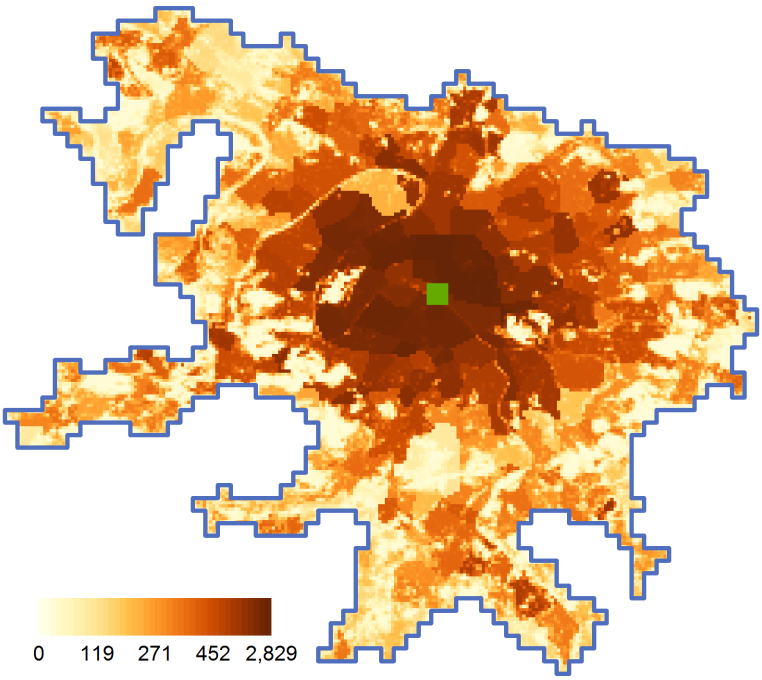

Note: The blue polygons show the delineation I use for Paris. The boundaries are taken from Florczyk et al. (2019). The green squares show my preferred city center definition. The coordinates are taken from OpenStreetMap and coincide with the Google Maps city center. The black dots in Subfigure (a) show the distribution of the 106,684 Airbnb properties in the city that were active during the study period. The raster in Subfigure (b) shows population counts in 250 meter × 250 meter grid cells, taken from Schiavina et al. (2019).

Figure A2: Example of city definition: Den Haag and Rotterdam

(a) Distribution of Airbnb properties

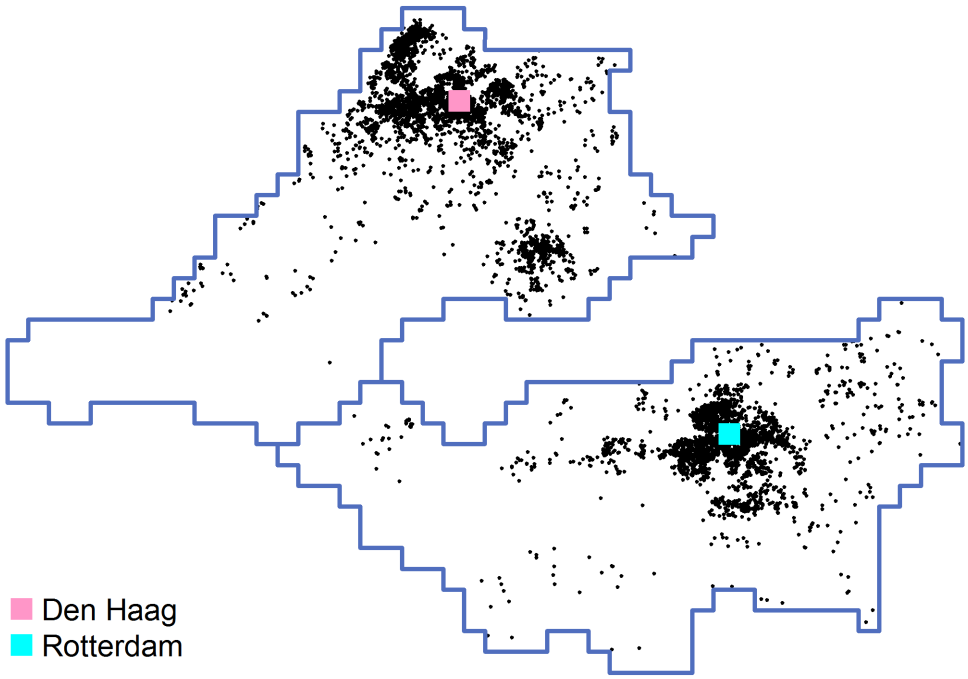

(b) Population count in 250 × 250 meter grid cells

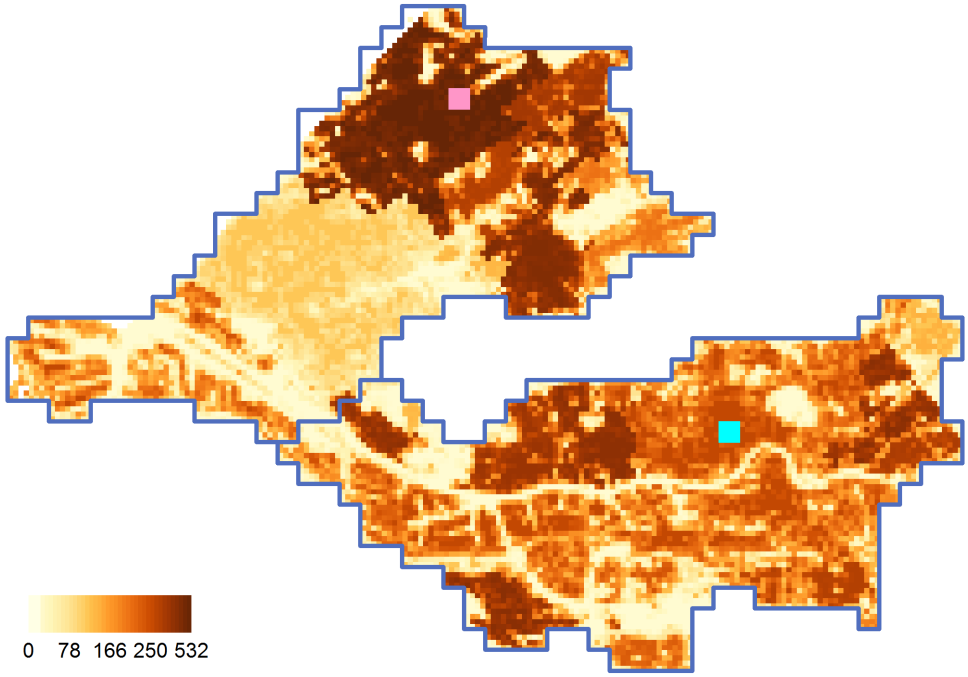

Note: The blue polygons show the delineation I use for Den Haag and Rotterdam. I split the boundaries of a combined urban center from Florczyk et al. (2019) by using the procedure described in Point 6 above. The pink and light blue squares show my preferred city center definitions for Rotterdam and Den Haag respectively. The coordinates are taken from OpenStreetMap. They coincide with the Google Maps city center for Rotterdam, but not for Den Haag. The black dots in Subfigure (a) show the distribution of the 3,282 Airbnb properties in Den Haag and the 2,638 Airbnb properties in Rotterdam that were active during the study period. The raster in Subfigure (b) shows population counts in 250 meter × 250 meter grid cells, taken from Schiavina et al. (2019).

Figure A3: Example of city definition: Rosario

(a) Distribution of Airbnb properties

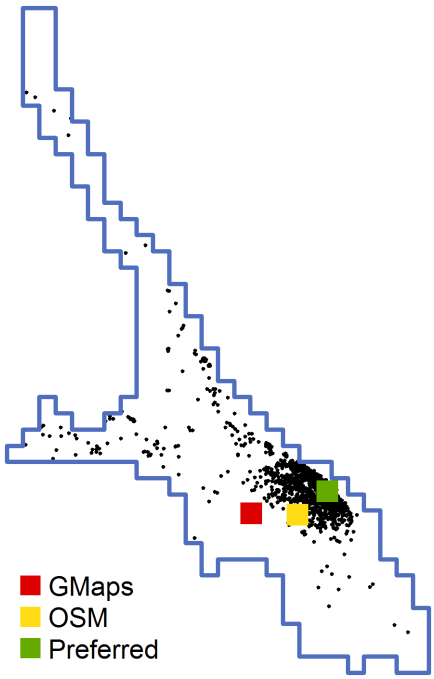

(b) Population count in 250 × 250 meter grid cells

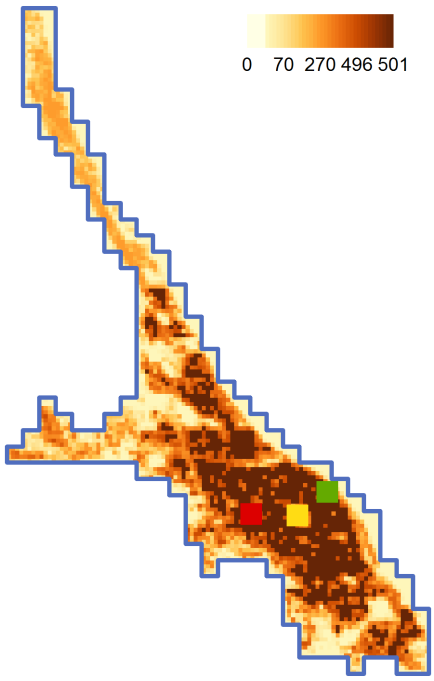

Note: The blue polygons show the delineation I use for Rosario. The boundaries are taken from Florczyk et al. (2019). The green squares show my preferred city center definition, the yellow squares show the definition according to OpenStreetMap, and the red squares show the definition according to Google Maps. The black dots in Subfigure (a) show the distribution of the 1,307 Airbnb properties in the city that were active during the study period. The raster in Subfigure (b) shows population counts in 250 meter × 250 meter grid cells, taken from Schiavina et al. (2019).

Figure A4: Number of cities included in the sample

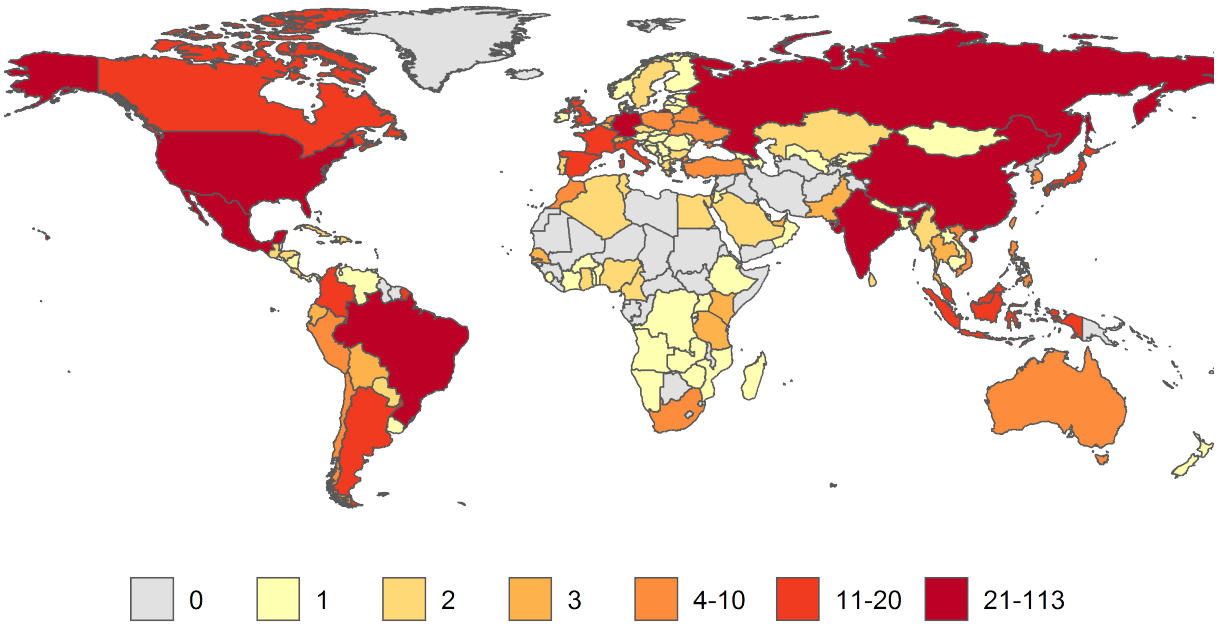

Note: This figure shows the geographic distribution of the 734 cities in the sample. To be included, a city must have at least 300,000 inhabitants and at least 100 Airbnbs that have been rented at least once over the sample period. The exact counts that are not visible from the map are: China (113 cities), United States (70), Brazil, Russia (both 44), Mexico (38), India (31), Germany (21), United Kingdom (20), Japan (17), Colombia (16), Spain (13), Argentina, France, Malaysia (all 12), Canada, Indonesia, Italy (all 11), Morocco, Philippines, Poland, South Korea (all 9), Peru, Taiwan, Ukraine, Vietnam (all 8), Australia, Belarus, South Africa, Turkey (all 6), Chile, Netherlands (both 5), and Isreal (4).

B Derivations of the implied transportation cost gradient

Start with the Alonso-Muth condition (from Alonso (1964) and Muth (1969); for the derivation of the condition see Duranton and Puga, 2015):

$$R'(x) = -\frac{t'(x)}{h(x)} \,, \tag{B.1}$$

where x denotes distance from the city center, $R(x)$ describes rents, $t(x)$ transportation costs, and $h(x)$ housing unit sizes. Assume that we are in a world where all buildings have the same height, while housing units can differ in size. In that case

$$D(x) = \frac{1}{h(x)} \,, \tag{B.2}$$

$$R'(x) = -t'(x)D(x) \,. \tag{B.3}$$

Estimating rents with regard to distance as log-log implicitly assumes the following functional form for $R(x)$

$$\ln\big(R(x)\big) = a + b\ln(x) \,, \tag{B.4}$$

$$\ln\big(R(x)\big) = a + \ln(x^b) \,, \tag{B.5}$$

$$\ln\big(R(x)\big) = \ln(Ax^b) \,, \tag{B.6}$$

$$R(x) = Ax^b \,, \qquad \text{where } A = \mathrm{e}^a \,. \tag{B.7}$$

It follows that

$$R'(x) = bAx^{b-1} \,. \tag{B.8}$$

If we assume that density is also best represented as log-log in distance, this implies

$$\ln\big(D(x)\big) = c + d\ln(x) \,, \tag{B.9}$$

$$\ln\big(D(x)\big) = c + \ln(x^d) \,, \tag{B.10}$$

$$\ln\big(D(x)\big) = \ln(Cx^d) \,, \tag{B.11}$$

$$D(x) = Cx^d \,, \qquad \text{where } C = \mathrm{e}^c \,. \tag{B.12}$$

Plugging B.8 and B.12 into B.3 yields

$$bAx^{b-1} = -t'(x)Cx^d \,, \tag{B.13}$$

$$t'(x) = -\frac{bA}{C}x^{b-d-1} \,. \tag{B.14}$$

Integrating with regard to x yields

$$t(x) = \int -\frac{bA}{C} x^{b-d-1} dx \ , \tag{B.15}$$

$$t(x) = -\frac{bA}{C} \int x^{b-d-1} dx \ , \tag{B.16}$$

$$t(x) = -\frac{bA}{C(b-d)} x^{b-d} + \text{constant} \ . \tag{B.17}$$

As we want $t(0) = 0$, the constant drops out and we are left with

$$t(x) = -\frac{bA}{C(b-d)} x^{b-d} \ . \tag{B.18}$$

C Additional regression results

Figure A5: Estimates of rental object-level hedonic regression

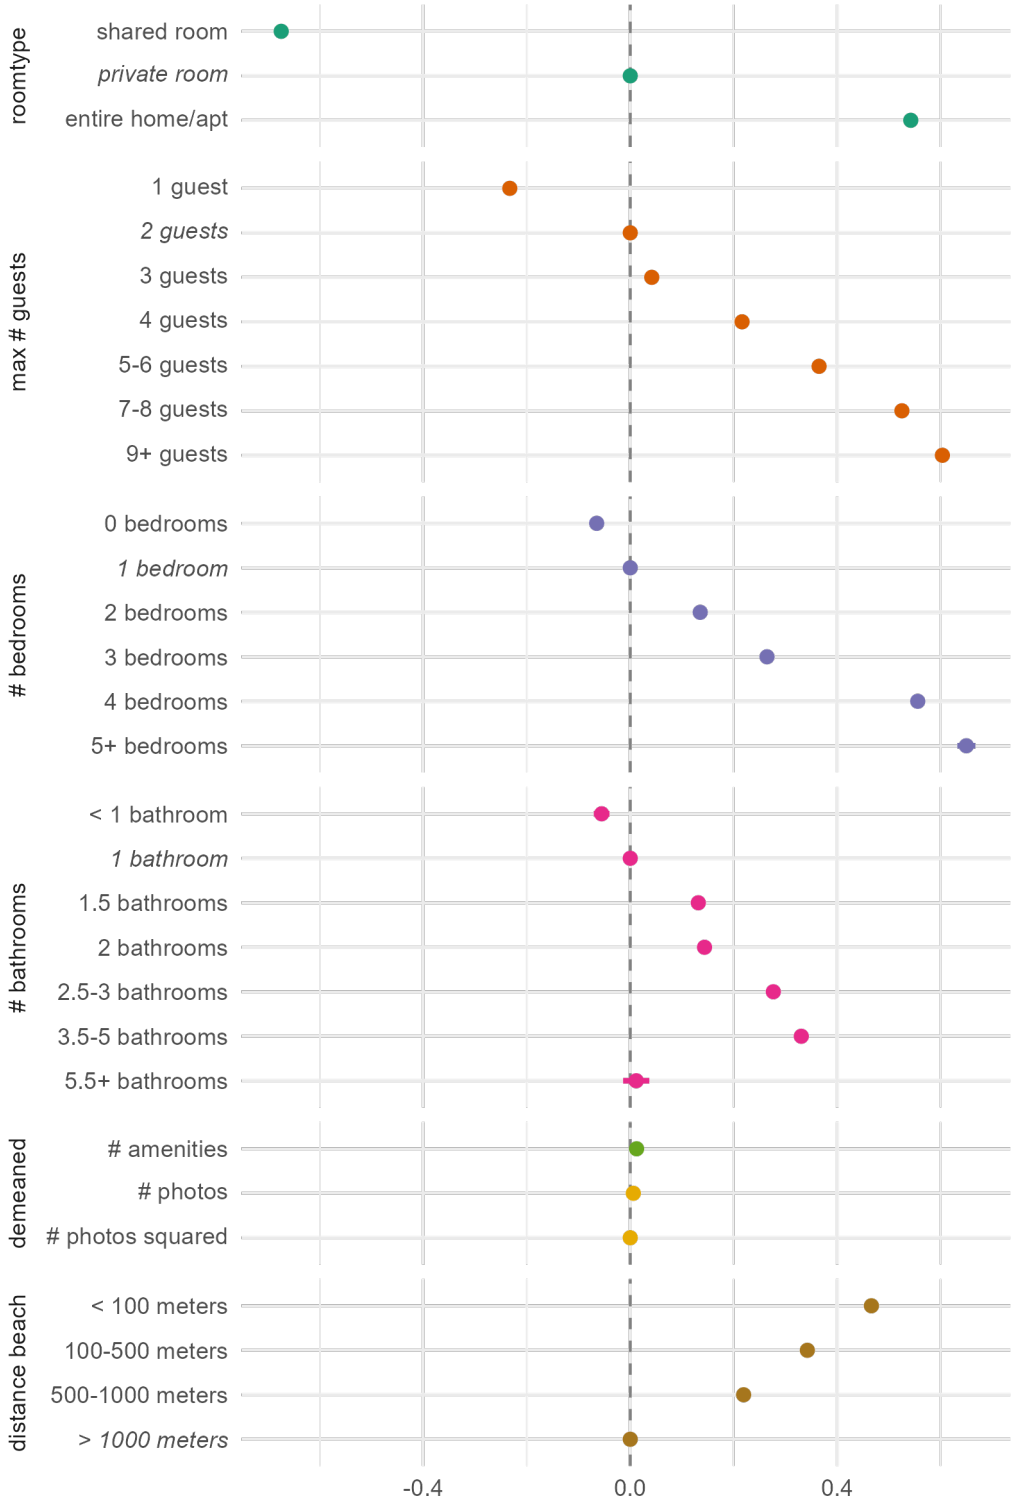

Note: This figure depicts the estimates of the first-stage hedonic regression of prices on property characteristics. The baseline categories are denoted in italic. Given the large number of 3,068,152 observations, most coefficients are so precisely estimated that the heteroscedasticy robust 99% confidence intervals are not visible.

Figure A6: Distribution of transportation cost gradients for cities with $b < 0$ and $d < 0$

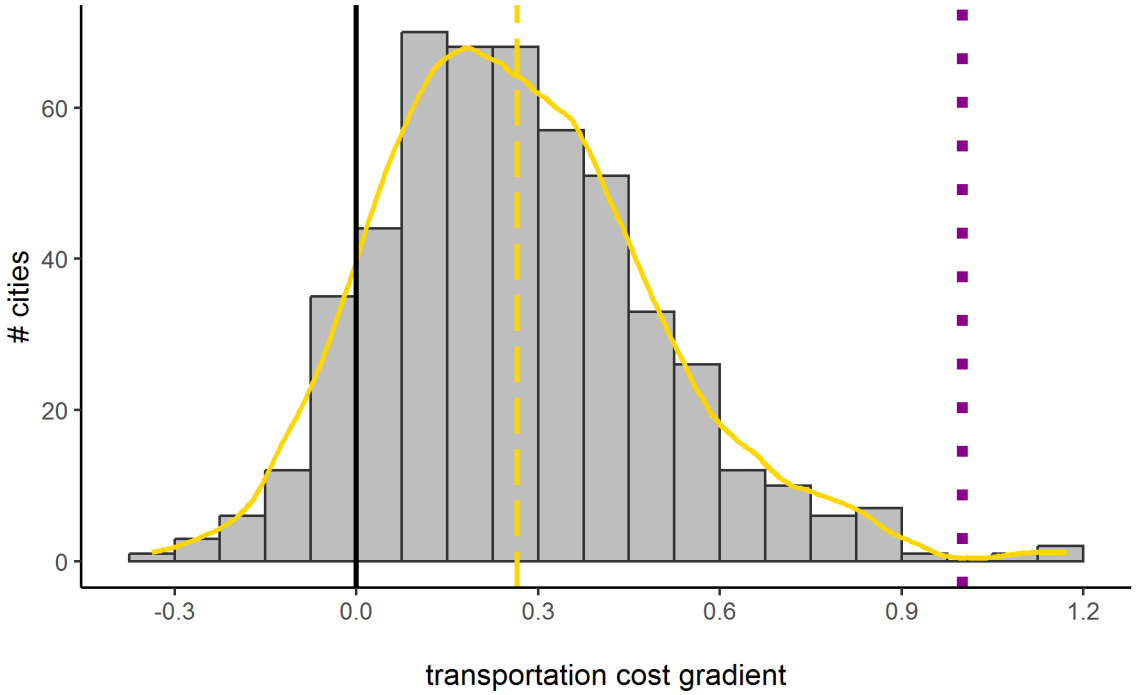

Note: Analogously to Figure 8, this figure shows the distribution of transportation cost gradients for all 734 cities in my sample. However, it only includes cities for which I estimate negative rent and density gradients, as predicted by the monocentric city model. The dashed line depicts the average transportation cost gradient, while the solid line shows the corresponding Epanechnikov kernel density estimate. The dotted line indicates an elasticity of transportation costs of one; equivalent to linear transportation costs.

Figure A7: Inferred transportation cost by distance,
relative to transportation costs for a 1 km trip

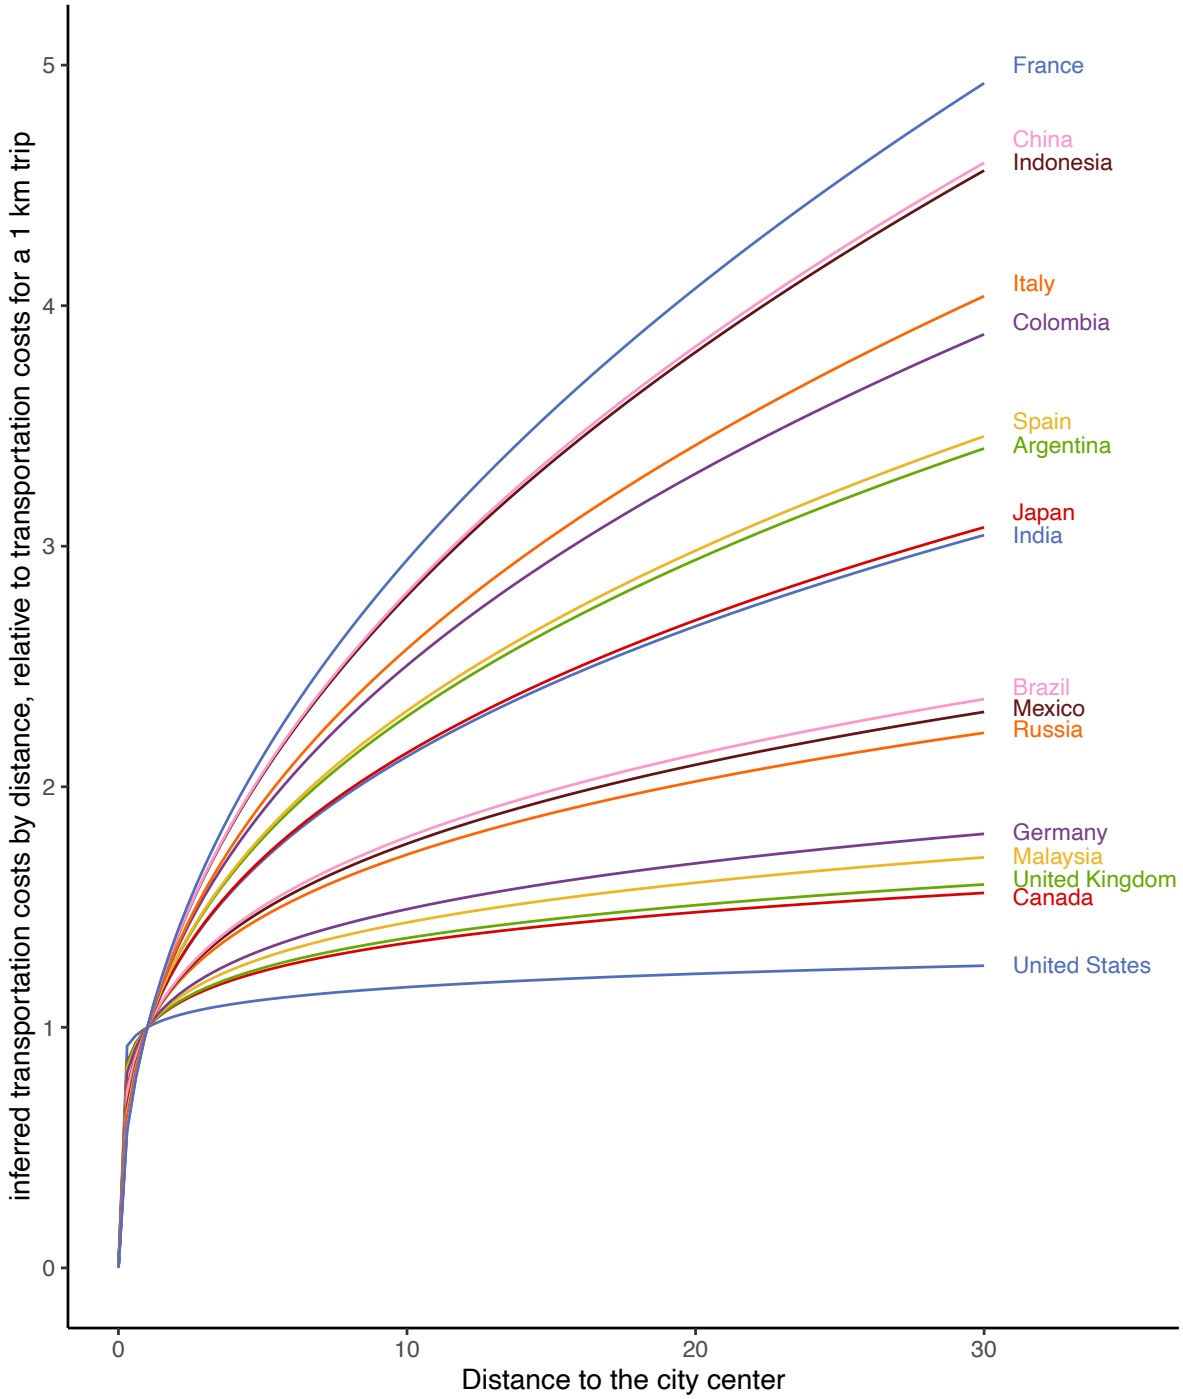

Note: This paper infers a transportation cost function $t(x) = \phi x^{\theta}$. In a monocentric city model setting in which rent and density are both assumed to follow a log-log relation with distance, θ is shown to be equal to the difference between the rent and the density gradient. I estimate both gradients and infer $\hat{\theta}$ for 734 cities worldwide. This figure depicts transportation cost functions for all countries with more than 10 cities in the sample. θ is set to the average $\hat{\theta}$ for each country, while ϕ is set to one. Importantly, the figure does not imply that France has higher transportation costs as the United States, as France might potentially have a much lower ϕ. Instead, the curves can be interpreted as the inferred transportation costs by distance, relative to the transportation costs for a 1 km trip. Thus, the figure suggets that the transportation costs of the first kilometer of transportation close to the city center account for a larger part of transportation costs in the United States than in France.

Chapter 2

How much more expensive is housing in larger cities? Worldwide evidence from Airbnb

How much more expensive is housing in larger cities? Worldwide evidence from Airbnb[1]

Bernhard Nöbauer[2]

University of Lausanne

September 2023

Abstract

Using a hedonic regression approach with data from 1.53 million Airbnb properties, I estimate the price of a representative short-term rental property at the center of 734 cities worldwide. The estimated rental prices provide an internationally standardized proxy for housing costs. Rental prices computed in this way are found to be highest in Amsterdam, London, New York, and San Francisco. I use these standardized rental price estimates to compute the elasticity of housing costs with respect to city size. My preferred specification shows an elasticity of 0.16, statistically significant at the 1% level. However, there is considerable geographic heterogeneity. Housing costs increase more strongly in city size in the euro area and India than elsewhere. In contrast, I find them to decrease in city size in Mexico. I offer suggestive evidence that crime might explain this unusual result.

[1] I thank Marius Brülhart for his guidance and countless hours of fruitful discussion. Moreover, I thank Dzhamilya Nigmatulina for helpful comments and Gilles Duranton for his inspiration. Moreover, I thank Laura Camarero Wislocka for her excellent help with assessing and defining city centers.

[2] Departement of Economics, Faculty of Business and Economics (HEC Lausanne), University Of Lausanne, 1015 Lausanne, Switzerland; bernhard.noebauer@bluewin.ch.

1 Introduction

A lot of empirical work confirms that wages are increasing in city size (for a survey see, for example, Combes and Gobillon, 2015). Bigger cities offer better opportunities to learn, share, and match, forces that are commonly summarized under agglomeration economies (Duranton and Puga, 2004). However, since our societies have not converged to live in a single gigantic city, there must be costs that make big cities less efficient or pleasant and that at least partially counteract agglomeration economies (Henderson, 1974). A particularly prominent example of such costs is housing costs, which are the focus of this paper.

The elasticity of housing costs with respect to city size measures how much more expensive housing becomes when city size increases, and estimates of that elasticity are surprisingly scarce. A prominent exception is Combes et al. (2018), who measure an elasticity of house prices with respect to city size of 0.21 for a sample of 277 urban areas in France. Given the French context, their estimates are mainly based on mid-sized cities, with an average urban area population of 166,020 and a median of 47,909 (p. 1565). However, we might expect urban costs to be disproportionally higher in the largest cities. Combes et al. (2018) show evidence of that by estimating a non-linear effect of city size, but these estimates are based on few observations at the upper end of the French city size distribution. I supplement their evidence using a worldwide sample of 734 cities with an average population of 2,100,936 and a median of 905,270 that has more to say about the housing cost premium of large and very large cities and goes beyond the context of a developed country.

Methodologically, I follow Combes et al. (2018) in measuring housing costs at the city center. This has the advantage that differences in transportation costs have a smaller influence on comparisons across cities, or no influence at all if we take the monocentric city model at face value. This model still guides a lot of research in urban economics and its assumptions are widely applied (for a survey on the model and its application see Duranton and Puga, 2015). In contrast, comparing city average real estate prices comes with the problem that the average property in a big city like Tokyo is further away from the center than the average property in a smaller city like Kagoshima, which implies higher transportation costs that confound the comparison.

For my analysis, I use data on short-term rental properties from Airbnb. Using these novel data allows me to extend the analysis to the global scale, based on an extensive set of variables that describe the properties in an internationally standardized way. This worldwide scope is hard to achieve with traditional data from national statistical offices or real estate platforms. As Airbnb hosts typically compete for the same housing units as long-term residents, across-city differences in nightly rates serve as a proxy for differences in long-term housing costs. I show, for the examples of France and the United States, that the city comparisons of housing costs estimated with Airbnb properties correspond to those estimated with long-term rental objects, albeit not perfectly.

I choose the 734 cities in the sample, their geographic boundaries, and their center points using transparent rules that I apply worldwide. Within these cities, I have data on 1.53 million properties

that were active between January 2018 and March 2019 and available or rented for at least 100 out of 365 days. I do this sample restriction to exclude apartments that only capitalize on peak price periods, using the Men's Fifa World Cup in June and July 2018 in Russia as a natural experiment to determine the cutoff value.

Using a hedonic regression with city fixed effects and city-specific distance gradients, I create a ranking of the 734 cities regarding the rental rate of a representative property at the city center. That property can be rented in its entirety by a maximum of two guests, has one bedroom and one bathroom, and shares the standard of an average Airbnb property in its city regarding all other characteristics. Given my methodology and data, I estimate Amsterdam, London, New York, and San Francisco to have the highest rental prices in the world. Caracas, Mandalay, Monteria, and Srinagar are at the other end of the ranking, with rental prices that are around 20 times lower. I run multiple robustness checks to confirm that the ranking is robust to changing underlying assumptions.

In the second stage, I regress the estimated rental prices on city size. For most specifications, I follow Combes et al. (2018) in using log population to measure city size and controlling for log area. This setup can be read as an unrestricted version of population density. I include country fixed effects when using the worldwide sample, so the coefficients are estimated from within-country variation. Moreover, I control for various city characteristics, including for the number of Airbnb properties per 100,000 inhabitants to control for the attractiveness of a city to tourists. An instrumental variable approach in which historical population sizes are used as an instrument serves as a robustness check. In my preferred specification, I estimate an elasticity of 0.161. This coefficient implies that a 10% higher population size is associated with housing costs that are 1.61% higher. The effect is statistically significant at the 1% level.

The literature provides a small number of related results. Ahlfeldt and Pietrostefani (2019) suggest an elasticity of rent with respect to population density of 0.15,[1] while Henderson (2002) estimates an elasticity of the rent to income ratio with respect to metro area size of 0.32. When not controlling for area, my results for the elasticity of housing costs with respect to city size almost precisely match those of Combes et al. (2018). They report an elasticity of 0.11, while my estimation yields an elasticity of 0.12. Combes et al. (2018) interpret this specification as the costs of unrestricted city size, while controlling for area corresponds to a city that is restricted from expanding outwards. The estimates suggest that the costs of unrestricted city size are very similar in our two contexts, with the difference between our main estimates coming exclusively from the area-restricted version. An intuitive explanation for this finding could the be stringent building height regulations in France (Jedwab et al., 2022). If a city is not allowed to expand outwards, constructing higher buildings is one of the remaining solutions to accommodate a larger population. The extent to which this

[1] Ahlfeldt and Pietrostefani (2019) is a meta-study that discusses the effects of density on multiple outcome variables. When I use density, instead of the more flexible specification of population and area, I obtain an estimate of 0.21, which is statistically significant at the 1% level.

solution is embraced will affect the increase in housing costs associated with a growing population.[2]

My work also expands the evidence on the geographical heterogeneity of the elasticity of housing costs with respect to city size. I compute separate regressions for the six countries with the highest number of cities in the sample (the United States, Russia, China, India, Brazil, and Mexico) and for the eurozone. The estimated elasticity is above the global average for the United States and Russia and is particularly high in the eurozone and India.[3] The estimate for the eurozone is within 0.04 percentage points from what Combes et al. (2018) estimate for comparably large French cities when using their non-linear specification. These findings point again towards an above average elasticity of housing costs with respect to city size for large European cities. While Chauvin et al. (2017) focus on agglomeration economies rather than urban costs, their work includes estimations of the elasticity of housing costs with respect to city size for the United States, Brazil, China, and India. My results are similar to theirs for the US, China, and Brazil, but they are very different for India, where Chauvin et al. (2017) do not find any effect of city size on housing rents. However, they do control for neither property nor city characteristics and they estimate the price of an average housing unit instead of a housing unit at the city center. When I apply their second-stage estimation strategy, I also find an elasticity that is indistinguishable from zero.

Being surrounded by many people might not always be beneficial. For Mexico, I estimate a statistically significantly negative coefficient for population and a statistically significantly positive coefficient for area. This finding implies that denser cities are cheaper in the Mexican context. I conjecture that crime might be a driver of this finding. The country is in the midst of a drug war (see, for example, Shirk and Wallman, 2015) and safety concerns are probably more important than elsewhere. I explore this hypothesis by adding an interaction term between log population and a city's homicide rate. My results show that Mexican cities with high homicide rates have a statistically significantly more negative elasticity of housing costs with respect to city size. I then proceed with the global sample and test the interaction between city size and an indicator for being among the 50 cities with the world's highest homicide rates. Cities in that group have a less positive elasticity of housing costs with respect to city size. The difference is statistically significant and quantitatively large, with an estimated elasticity that is more than 40% lower. The finding that crime lowers the elasticity of housing costs with respect to city size complements the evidence that large cities are more affected by crime (Glaeser and Sacerdote, 1999) and that crime negatively affects house prices (Pope and Pope, 2012).

The remainder of the paper is organized as follows: Section 2 describes the main data and discusses city definitions. Section 3 presents the first-stage hedonic regression and results in a ranking of cities by their estimated short-term rental price at the city center. In Section 4, I use these rental prices as an input for the second-stage regressions, and I present and discuss the corresponding

[2] French central cities often have many beautiful old buildings, and tearing them down is probably not an optimal solution. Glaeser (2011) discusses this nexus and possible ways forward using the example of Paris.

[3] The coefficients are statistically significant at the 5% level for the US, India, and the eurozone, but not for Russia.

estimation results. Section 5 concludes. Finally, the appendix contains the full city ranking and several auxiliary results.

2 Data and city definitions

This project relies on two main types of data: geolocalized data on Airbnb properties and spatially disaggregated population data. I combine the latter with data on city centers and the boundaries of urban areas to define 734 cities that I include in this study. All of these data are available in an internationally standardized way, which allows me to conduct the analysis on a global scale. This section will successively present the data on the Airbnb properties, the city definitions I use and the data on population sizes of the resulting cities.

Airbnb properties

The data on short-term rental properties from Airbnb come from AirDNA, a company specialized in "short-term rental data and analytics".[4] They contain close to all properties that were advertised on Airbnb at least once between 2018-01-01 and 2019-03-25.[5] By combining information about days for which properties are rented with information about the price for these days, AirDNA is able to estimate the prices actually paid by customers. For every property, I have information about the average daily price over the twelve months before the date on which a property was last web scraped from the Airbnb website. I also have the coordinates of the location for each property, even though some of them are scrambled within a short radius due to security concerns.[6] Moreover, the data contain a substantial number of covariates, from the number of bedrooms to the presence of a hairdryer. All of these variables are available in an internationally standardized way. Overall, I can match 3.07 million properties to the 734 cities in my sample, 1.53 million of which were available for rent or rented for at least 100 of the last 365 days before they were last scraped.[7]

Cities

It is not straightforward to find a definition of where a city ends and where its center is located. The problem becomes especially complicated if the definition is supposed to work well for very different

[4] `https://airdna.co`, last accessed: 2023-01-16.

[5] To the best of my knowledge, AirDNA web scraped every single property from Airbnb once every three days over this period. This implies that a small number of properties that appeared only briefly and were immediately removed or rented might not be part of the dataset.

[6] Airbnb recommends that hosts indicate their precise address. However, hosts are free to choose whether they prefer a precise pin to be shown at the address of their property, or a circle that indicates the approximate location in a close radius (`https://www.airbnb.com/resources/hosting-homes/a/setting-expectations-with-an-accurate-location-491?_set_bev_on_new_domain=1686948204_M2FkMTUxYWQzNzNi&locale=en`, last accessed: 2023-06-16). The current maximal deviation from the true location of the property is indicated as 800 meters. When I obtained the data in 2019, AirDNA suggested an even smaller maximal deviation of 500 meters.

[7] The raw dataset contains 9,419,495 observations. However, 2,354,445 of these properties were never reserved. I have to drop another 1,917 observations because their coordinates are missing. Afterwards, I can spatially join 3,093,755 properties with my city polygons. An additional 25,603 observations drop out because of missing covariates (or a missing price in one case). In the end, 3,068,152 entries remain.

countries. Here, I explain the main decisions I make to come to a definition that I deem suitable for the empirical exercise I conduct.[8]

I start with the open collaboration database platform OpenStreetMap that relies on crowd intelligence. This website asks users to place so-called city tags "at the center of the city, like the central square, a central administrative or religious building or a central road junction".[9] These geolocalized tags are my first candidates for both cities' locations and city centers. As a second step, I spatially join the city tags to all urban center polygons of the Global Human Settlement Layer project (Florczyk et al., 2019) with a population size of at least 300,000.[10] In some cases, urban center polygons contain multiple city tags. Using the city population counts from OpenStreetMap, I retain all city tags that are associated with a population count that amounts to at least 40% of the highest population count in an urban center.[11]

Figure 1: 734 cities in the sample

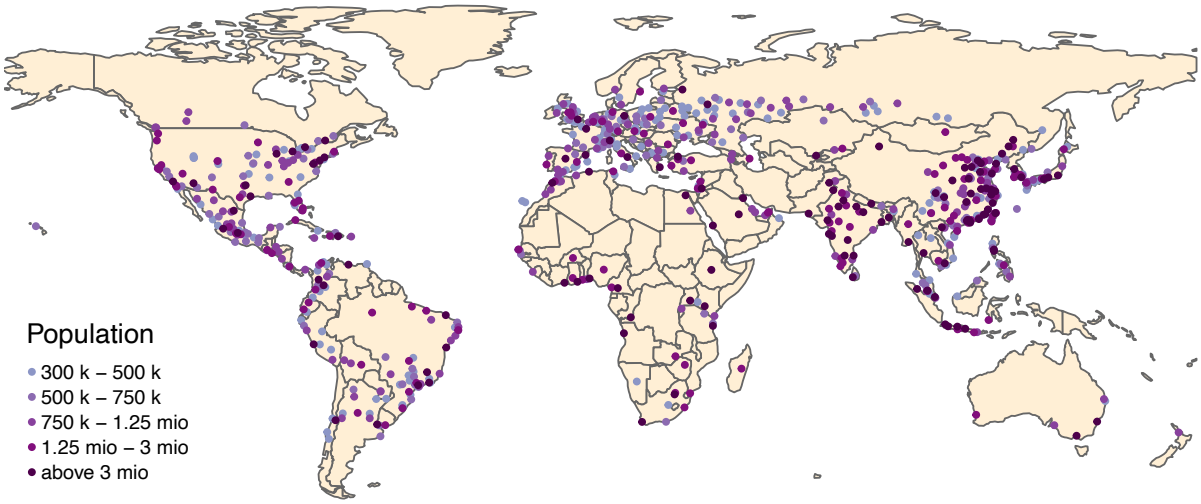

Note: The dots in this figure show the geographic distribution of the 734 cities in my sample. Their colors refer to the city's population size, with larger cities represented in darker shades. To be included in the sample, a city must have had a population of at least 300,000 inhabitants in 2015 and at least 100 Airbnb properties that were active between January 2018 and March 2019.

I then manually check all remaining tags using satellite and street view images from GoogleMaps

[8] For more details, refer to Nöbauer (2023), where I use the same city definitions and delineations and describe them more extensively.

[9] `https://wiki.openstreetmap.org/wiki/Tag:place%3Dcity`, last accessed: 2023-01-13. The website also includes information about many different kinds of geographic tags like motorways, restaurants, or playgrounds. As of 2023-06-12, the website contains 2.951 billion tags in total (`https://taginfo.openstreetmap.org/reports/database_statistics`).

[10] The basis of the urban centers of Florczyk et al. (2019) are contiguous 1km × 1km grid cells with an estimated population of at least 1,500 or a built-up area of at least 50%. Their definition results in some cities being very broad and containing multiple well-known cities, for example Oakland/San Francisco/San José or Kobe/Kyoto/Osaka. In some of these cases setting one city center for the whole urban area would be very tricky. I therefore decide against simply adopting their definition. I believe that combining the urban centers from Florczyk et al. (2019) with data on cities from OpenStreetMap results in a set of cities that is better suited for this analysis.

[11] I complement missing population counts on OpenStreetMap with information from Wikipedia.

and assess whether they are an appropriate choice for the city center. Whenever this is not the case, I resort to center coordinates from Google Maps. If they also describe a point that visually does not constitute a suitable center, I provide my own best guess. Once the city centers are determined, I split urban centers with multiple remaining city tags that are more than 7 km apart, so that it is likely that they constitute two distinct cities.[12] In a final step, I recompute the population size and the number of Airbnb properties within each city and retain those with at least 300,000 inhabitants and 100 properties.[13] Figure 1 shows the geographic distribution of the resulting 734 cities.

Population counts

The population counts also come from the Global Human Settlement project, more precisely from the GHS-POP file (Schiavina et al., 2019). They take population data from administrative sources at the smallest availbale scale. They then disaggregate these data to 1 km $\times$ 1 km grid cells using the proportion of buildings and other artificial structures, detected from day-light satellite images, with machine-learning.[14] They apply the same procedure to satellite images from 2015 and 1975, which ensures a certain level of intertemporal comparability that is beneficial to my instrumental variable approach.[15]

3 First-stage regressions

When assessing the effect of city size on real-estate prices, it is important to use housing units that are as comparable as possible across cities. Even with standardized data from a single source, simply computing the average price of units for each city is insufficient. There are two main reasons for this: First, the size and quality of housing units vary non-randomly. For example, the average apartment in Paris has fewer bedrooms than the average apartment in Toulon, while the average one-bedroom apartment in Ho Chi Minh City has more amenities than the average one-bedroom apartment in Can Tho. Second, differences in the geographical expanse of cities imply differences in accessibility and transportation costs. The average apartment in Buenos Aires is much farther away from the city center than the average apartment in Salta, and its inhabitants might spend considerably more time commuting for work and leisure activities than their counterparts in Salta.

To address the first issue, I estimate a hedonic regression. Apart from a separate intercept for each

[12] To split the cities, I use a rule described by Akbar et al. (2021): Border points X are assigned such that $\text{dist}(X,A)/\text{dist}(X,B) = (\text{Pop A}/\text{Pop B})^{\frac{0.57}{2}}$ where $\text{dist}(X,A)$ denotes the distance of a grid point X to the center of city A and $\text{dist}(X,B)$ denotes the distance to the center of city B. I reassign enclaves in repeated iterations until there are no city parts left that do not contain a center. I split cities that span across two countries at the border, without reassigning enclaves.

[13] To assess the population size I use the population data presented below. As discribed above, the cutoff of 100 Airbnbs refers to properties that have been rented at least once.

[14] They also offer a 100 m $\times$ 100 m resolution. However, I keep the 1 km $\times$ 1 km grid structure of the GHSL urban centers for my cities, so I would not gain anything by using the better resolution.

[15] As discussed in Section 4.1, 40 years are hardly enough for a credible identification of the instrument. Nevertheless, it is progress to have intertemporally comparable population data on a global scale. I therefore present the results of an IV specficiation, while cautioning that they should not be interpreted as more than a robustness check.

city, which is the variable of interest in this first stage, I control for numerous characteristics for each property, particularly for the number of bedrooms, bathrooms, and the maximum number of guests allowed.[16] I also include indicators for whether guests have the entire apartment for themselves or have to share the apartment or even their room. As the effect of these core characteristics may well be nonlinear, I allow for a flexible functional form by including them as categorical variables. The left column of Figure A1 shows the respective categories and their distribution in the data.

The same column also shows the distribution of other variables for which I control. The number of photos serves as a proxy for how much effort is put into creating the profile on Airbnb. For this variable, I also include a squared term, as I expect the marginal effect of additional photos to be diminishing and potentially even negative at a very high number of pictures. The number of properties a host has on the platform controls for the fact that certain hosts offer multiple properties. The final row of Figure A1 shows an indicator for whether a property is within 500 meters of an ocean or big lake.[17] Furthermore, I control for 43 amenities, examples of which include the presence of a tv, a hairdryer, or a first aid kit, as well as the availability of breakfast or free parking. Figure A2 displays the list of amenities, with the fraction of properties in which they are available in brackets.[18]

The variables in this second group are included either as indicators or modeled using a linear or quadratic functional form. Moreover, I demean them within each city. To see why this improves the estimation, consider the amenity "heating". Without demeaning, there is a selection effect. Most properties without heating are located closer to the equator. They are not necessarily cheaper because of the lack of heating, which is unnecessary in the warmest climate zones. However, they are often located in countries with lower overall price levels.[19] Including this variable without demeaning would therefore result in an overestimation of the effect of heating by absorbing part of the city-fixed effects. As this first-stage regression aims to estimate the city-fixed effects as precisely as possible, demeaning helps avoid these biases.

To address the second issue, I follow Combes et al. (2018) in estimating the price of a property at the city center rather than the price of an average property in a city. The economic intuition

[16] Unfortunately, I do not have data about the square meter size of an apartment. However, customers usually do not have access to this information either. It is only available to them if the host explicitly puts it in the property description or if they have stayed there before. In all other cases, customers cannot consider it for their decision-making, and I, therefore, expect its influence on prices to be limited.

[17] This is the only variable that is not directly visible on Airbnb. Instead, the customers can infer it from a map provided on the website, although Airbnb sometimes scrambles the coordinates to some limited extent (500m at the very most) for security concerns. Moreover, the hosts seem to have a clear incentive to indicate a location close to a coast or beach in the description and the photos. To construct these indicators, I measure the air-line distance from a property to the closest ocean, sea, or big lake (at least 80km^2). To determine the location of waters, I use ESRI's "World Water Bodies" layer (`https://arcgis.com/home/item.html?id=e750071279bf450cbd510454a80f2e63`, downloaded on 2023-10-10) and the HydroLAKES data from `https://hydrosheds.org/products/hydrolakes` (downloaded on 2023-01-01).

[18] I only include amenities that are present in at least 1% of properties. The data include another 34 amenities available in very few apartments.

[19] To some extent, this can also be the case within an individual country. For example, in Italy, heating will be more of a necessity in the northern part of the country, which is also the wealthier part of the country.

for this builds on two of the most well-known models in urban economics: The Rosen-Roback model (Rosen, 1979; Roback, 1982), which describes choices between cities, and the monocentric city model (Alonso, 1964; Mills, 1967; Muth, 1969), which describes choices within cities. The monocentric city model features households that work in the city center for a given wage, bear transportation costs for their commute, consume housing, and a composite good. In equilibrium, the unit cost of housing is more expensive closer to the city center, as people are willing to pay higher prices to avoid commuting costs. Ex-ante homogeneous agents can end up with different bundles of a location, housing consumption, and the composite good, with all bundles yielding the same utility. As long as this equalized within-city utility is given, it does not matter which bundle is taken for the comparison of agents across cities. It is convenient to make the comparison in the city center, where transportation costs are zero according to the model's assumptions. This choice, in turn, facilitates the comparisons between cities that underpin the utility equalization across cities in the Rosen-Roback model. In the model's equilibrium, wage differences and amenities counterbalance differences in housing costs. Measuring the housing costs at the city center implies that transportation costs can be left out of the comparison.[20] While the present paper is exclusively concerned with estimating the housing cost aspect of this comparison, it is important to bear this bigger picture in mind.

Empirically, I implement the measurement at the city center by estimating both an intercept μ_c and a distance gradient β_c for each city c. I add +1 to the distance to the city center to be able to interpret a distance of $\ln(1) = 0$ as the city center. My preferred first-stage regression has the form

$$\ln(\text{price})_{ic} = \mu_c + \beta_c \ln(\text{distance}+1)_{ic} + \boldsymbol{\gamma}\mathbf{X}_{ic} + \boldsymbol{\delta}(\mathbf{Z}_{ic} - \bar{\mathbf{Z}}_c) + \varepsilon_{ic} \,, \tag{1}$$

where X_{ic} denotes a set of core categorical variables for the type of the listing and the number of bedrooms, bathrooms, and the maximum number of guests allowed. The baseline categories are the respective modes (see Table A1). $\mathbf{Z}_{ic}$ denotes the second set of variables and amenities that are included with a specified functional form and demeaned by city.

A city fixed effect μ_c, therefore, has the interpretation of the log price of an apartment in city c that is located at the city center, rented out in its entirety to a maximal number of two guests, has one bedroom and one bathroom, and characteristics that match the city average for all variables in $\mathbf{Z}_{ic}$. Taking the exponential of the city fixed effect yields the USD price of this representative property.

3.1 Excluding apartments available only during price spikes

The founders of Airbnb got the idea for their business when participants of a conference in San Francisco struggled to find available hotel rooms (Gallagher, 2017). Some long-term tenants also

[20] In reality, people face transportation costs even if they live at the very center of a city, and these transportation costs vary across cities. However, I consider it probable that the comparison at the city center minimizes both the level of transportation costs and the differences in transportation costs between cities.

rent out their apartments while they are on vacation. In these cases, Airbnb can contribute to a more efficient capacity utilization of living space rather than merely displacing one kind of occupant with another. However, these cases also threaten the validity of using Airbnb data for my study. Major events can lead to a temporary surge in price. Examples include events with changing venues, such as the Super Bowl, but also annually recurring events like Art Basel. One might argue that such events contribute to the general attractiveness of a city and that an increase in price is, therefore, justified. However, if there is also a surge in the properties offered on Airbnb to take advantage of the temporarily higher prices, these marginal properties will have an average nightly rate that is much higher than what could be charged on a yearly basis. These properties will therefore bias the estimated prices for the concerned cities upwards. Before constructing a ranking of cities by their price level, I will try to mitigate that problem by excluding properties that are only on the market for a short time to capitalize on exceptionally high prices.

I do this by imposing a minimum number of nights in which a property is either reserved, or free and available for reservation. However, it is not evident how to choose a suitable cutoff. I use the Men's FIFA World Cup that took place in Russia from June 14th to July 15th, 2018, as a natural experiment. My sample contains 44 Russian cities, nine of which hosted games during the tournament.[21]

I run the first-stage regression on multiple subsets with increasingly strict cutoffs for the minimum number of nights on the market. The first set includes all properties in my sample of cities. The second subset only includes properties reserved or available for at least 25 out of 365 nights. I then proceed in steps of 25 nights, eventually reaching the strict requirement of 200 nights. For each regression and all Russian cities in the sample, I estimate the price for a representative property in the city center as defined above.

Figure 2 shows the results of this exercise. It depicts Russian cities that hosted World Cup games in blue and cities that did not host world cup games in yellow. Host cities are inherently different. Prices in these cities are higher even for properties on the market for most nights. This regularity makes intuitive sense, as games are usually played in larger cities that can provide the required infrastructure. However, more relevantly, the price gap between host and non-host cities decreases in the number of nights on the market. It is most prominent for the whole sample without restrictions and then declines monotonically for most pairs of cities. Depending on the city, removing properties that were only on the market for less than 25, 50, or 75 nights leads to a substantive drop in the estimated price. The decline then fades out, with a modest change associated with removing properties that were reserved or available between 75 and 100 nights. Removing properties that were on the market beyond 100 nights does not change the estimated prices in any significant way, with the blue lines becoming essentially horizontal. With at least 100 nights on the market, the remaining properties will hardly be inhabited by ordinary long-term tenants capitalizing on major

[21] In total, the world cup was played in 12 stadiums, with Moscow featuring two venues. However, Saransk and Sochi do not meet the cutoff of 300,000 inhabitants to be included in the sample.

Figure 2: Price of a representative apartment at the city center: Russia

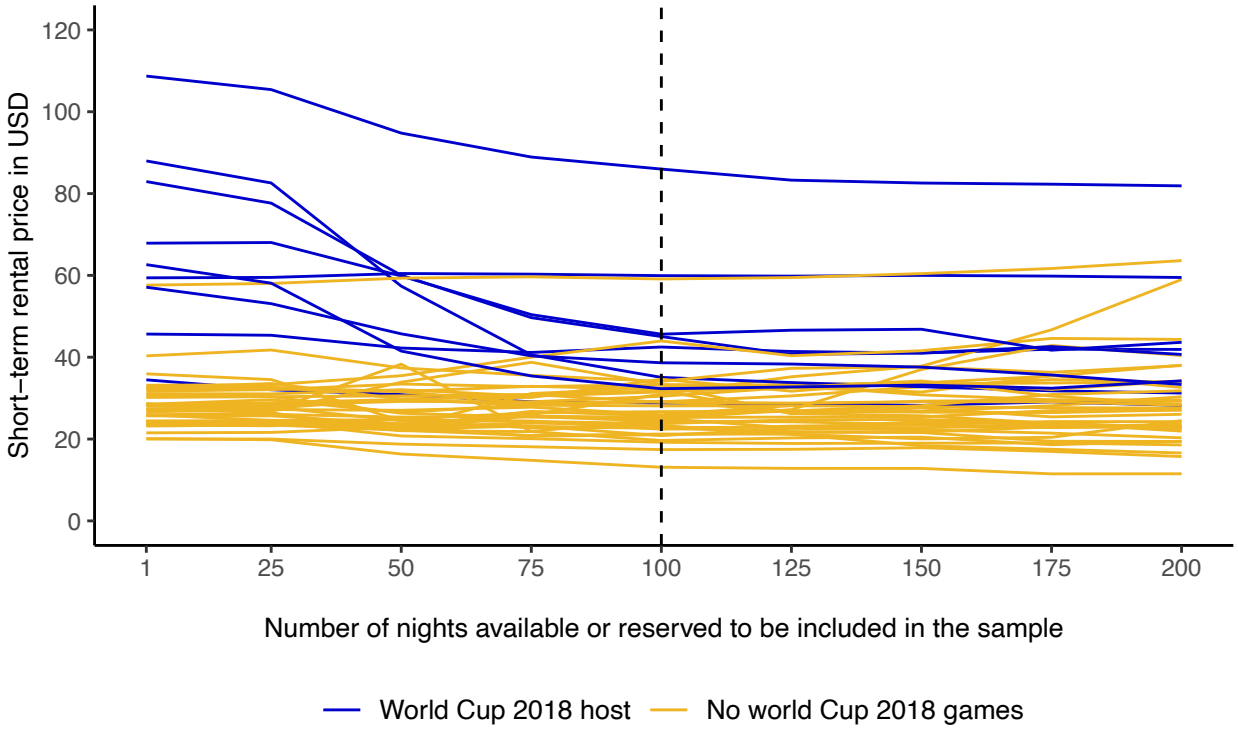

Note: The figure shows nightly US Dollar prices of a representative short-term rental property at the city center for the 44 Russian cities in my sample. I estimate the prices by taking the exponential of city-fixed effects that are the output of 9 different hedonic regressions. These hedonic regressions differ in the subset of Airbnb properties they consider. The 44 coefficients at the right end of the x-axis are based only on properties that were available or rented on at least 200 of the last 365 nights before a property was last scraped. When moving further left on the x-axis, weaker cut-off values apply. Cities that hosted games during the Men's Fifa World Cup 2018 are shown in blue, while cities that did not are depicted in yellow.

price surges.

Non-host cities show no general trend along the whole spectrum of subsets, with the estimates becoming slightly more dispersed as the considered properties get scarcer. While these cities certainly also have varying demands over the year, only major events seem to lead to a notable rise in Airbnb properties supplied that can explain the pattern of prices. It is reassuring that one non-host city also features higher prices over the year while still displaying a stable estimated price across the different subsets. This city is Vladivostok, which was disregarded as a venue to have shorter travel distances (FIFA, 2010).

As a consequence of this analysis, I limit the data to properties that were reserved or available at least 100 out of 365 days. Due to this restriction, I lose about half of the properties, leaving me with slightly more than 1.53 million observations.

3.2 First-stage results

I then continue to estimate equation (1). Table A1 lists all 734 cities by their nightly short-term rental rate of a representative property at the city center. As defined above, that property can be rented in its entirety by a maximum number of two guests, has one bedroom and one bathroom, and corresponds to the average within its city regarding all other characteristics. Given the data and methodology I use, the most expensive city is Amsterdam, with a nightly rate of 252 USD, followed by San Francisco (243 USD), London (231 USD), and New York (225 USD). My second-stage regressions are based on (the logs of) these prices.

While these cities on top of the list are all infamous for high housing prices, they are also major tourist destinations. Perhaps as a consequence of this combination, all four cities introduced some regulation regarding Airbnb properties early on (see, for example, von Briel and Dolnicar, 2021). I cannot rule out that differences in the strictness of these regulations influence prices for such illustrious cities. However, given this study's large number of cities, I do not expect this to be a significant issue for my second-stage analysis. To control for the exposure to tourism, I include the number of Airbnbs per 1,000 inhabitants as a control variable in the second stage.

The right column of Figure A1 shows the estimated coefficients of the control variables for the first-stage hedonic regression. The number of bedrooms, bathrooms, and maximum allowed guests all increase the price of a property monotonically, although the differences between the individual coefficients vary. For example, a host can charge substantially more if she allows two guests instead of one. Hosting three instead of two guests increases prices by much less. Compared to the other coefficients, the type of the listing has a large effect on prices, with shared rooms coming with a particular markdown.

More photos in the ad correlate with higher prices, with a slightly diminishing marginal return. Apartments offered by hosts with several properties on the platform are more expensive. These hosts might learn to optimize traveler experience when spending a lot of time on the platform, allowing them to charge higher prices. They might also be able to charge more if they have market power on particular submarkets. Finally, properties close to an ocean or big lake are more expensive.

Figure A2 assesses the effects of various amenities. The amenities most positively related to price are air conditioners, pools, and TVs. There are also amenities that are negatively correlated with price. The heterogeneity of these cases suggests there are several different explanations for this. Examples include situations where the necessity of an amenity points towards an inconvenience, like a lock on the bedroom door or room darkening shades, as well as amenities that signal that an apartment is not optimized for travelers, like children's books and toys or a washer.

3.3 Robustness checks

I perform extensive checks to assess the robustness of my first-stage results. For each version, I recalculate a ranking with the estimated prices at the city center. Table A2 reports the correlations

between the prices estimated with the different specifications, while Table A3 shows corresponding Spearman's rank correlations.

Specification A denotes the baseline version described above. Specification B restricts the sample to the category "entire home/apartment," dropping listings classified as "private room" (29% of all listings) and "shared room" (2%). Entire apartments are certainly what first comes to mind when thinking about the long-term rental market and might therefore appear to provide the closest correspondence. However, people also share apartments or rooms for extended periods, a prominent example being university students.

The data on Airbnbs also contain information about customer ratings, albeit for a reduced number of properties (74% of all properties in my sample of cities, but 90% of properties that were reserved or available at least 100 out of 365 days). Specification C controls for the demeaned ratings in the following categories: accuracy, check-in, cleanliness, communication, and value. I abstain from including the location rating as it might interfere with the distance gradients. Specification D controls for all amenities in the data, even if they are available only in very few properties.

There is a slight subtlety concerning the control for a location within 500m of an ocean or big lake: I demean all variables other than the number of bedrooms, bathrooms, the maximal number of guests, and the listing type by city. This implies that the representative property located at the city center has the characteristics of an average property in the city. However, when it comes to proximity to a large water body, the city center is either located close to a big lake or ocean, or it is not. In that sense, controlling for demeaned water proximity implies a somewhat flawed interpretation. On the other hand, omitting it means ignoring a factor that considerably impacts real estate prices while being correlated with proximity to the city center. Therefore I report both, with specification E omitting demeaned water distance.

Specifications F and G revisit the exclusion of properties that have been on the market for less than 100 of 365 days. Specification F includes all properties reserved at least once and thus having a revealed price. Including them approximately doubles the number of properties I can use to compute the first-stage regressions. However, it can lead to overestimating prices for cities with important events during a limited number of days. In contrast, specification G applies a more stringent requirement of 125 days on the market, resulting in a drop of another 240,000 properties. A few prices seem unrealistically high and are most likely erroneous. Therefore, my main specification winsorizes prices to each country's 0.01 and 0.99 percentiles. Specification H uses the prices at face value without winsorizing.

Specifications I and J refer to the choice of the city center. As described in Section 2 and in Nöbauer (2023), my city centers are mainly based on city tags set on OpenStreetMap using crowd intelligence. Together with a research assistant I evaluated all of these centers. Whenever they do not withstand a visual assessment, I continue with Google Maps city coordinates. If these are also suboptimal, I propose my own best guess. In contrast, specification I uses the coordinates from OpenStreetMap

for all cities, while specification J uses the coordinates from Google Maps. Moreover, specification K uses ln(distance) instead of ln(distance + 1) to compute the distance gradients.

Specification L introduces a new way of demeaning. It demeans the same variables as the baseline specification. However, instead of demeaning them by city, I split the properties in each city into two subsets according to their air-line distance to the city center. I then demean the variables by city halves. This procedure is less prone to confound the effects of amenities with the distance gradients. For example, the amenity street parking negatively affects a property's price. This amenities suggests the absence of a garage or another secured parking facility. However, very central properties might not have any parking possibility and may yet be highly attractive. Therefore, part of the negative effect of street parking might be because it is a proxy for non-central locations. This problem is alleviated by demeaning within the groups of more central and less central properties in each city.

Overall, the results of the robustness tests are reassuring. The median price correlation across specifications A to L is 0.98, and the median rank correlation is 0.95.

Finally, specification M computes a ranking based on estimated average rental prices in the city, instead of estimated rental prices at the city center. It is based on the same first-stage hedonic regression as the other specifications, but it does not include distance gradients. This specification exhibits substantially lower correlations with the other specifications. However, the correlations are still 0.92 (prices) and 0.89 (rents), which might not be surprising given the large international differences in price levels.

3.4 Comparison with longterm rental data

In principle, people offering Airbnb properties compete for the same apartments as long-term renters. Living space is the primary input for the service offered by Airbnb. There is no apparent reason why other inputs like furniture or labor conducted by cleaners should vary differently between locations for the two markets. Therefore, in a market economy, we can expect Airbnb prices to be high in places with high long-term rentals and vice versa.[22] However, long-term rentals are substantially regulated in some countries, especially concerning existing tenants. Therefore, the results in this paper reflect the market for new long-term rentals more closely, as they are usually less regulated.

As a further validity test, I compare my estimates to center prices estimated using long-term rental data. I do this for France and the United States; a choice driven by data availability. Unfortunately, I do not have access to property-level long-term rental data, so I rely on aggregated data on a granular geographic dimension.

For France, I work with *la carte des loyers*.[23] This map is provided by the French government

[22] The period covered by my data ensures that the estimates are not influenced by the Covid pandemic with all its implications, which were very different for the two sectors.

[23] The data can be found on `https://www.data.gouv.fr/fr/datasets/carte-des-loyers-indicateurs-de-loyers-`

and is based on 7 million real estate ads posted between 2018 and 2022 on `seloger.com` and `leboncoin.fr`. It displays prices estimated by a hedonic regression for 34,980 French communes and arrondissements.[24] 451 of these geographical units are located in one of the 12 (functional) French cities in my study.[25] I regress the log of the rents from *la carte des loyers* on a city intercept and a city distance gradient for each city in my sample. Panel A of Figure A3 presents the results. The axes depict the city intercepts, which can be interpreted as the log prices for representative properties at the city center, once estimated using data from Airbnbs and once from *la carte des loyers*. The correlation is relatively high. Both sources also consistently estimate the extent to which Paris is an outlier in the French context; a regularity that is also found by Combes et al. (2018).

For the United States, I use data from the American Community Survey. The data cover 2015-2019 and are spatially disaggregated at the block group level.[26] I regress the log of median rents on a city intercept and city gradient for each city. Unlike above, the rents are summary statistics from survey responses rather than the result of a hedonic regression. Therefore, I control for a list of covariates linked to real estate at the block group level.[27] Similar to above, Panel B of Figure A3 compares the city intercepts estimated using long-term and short-term rental data for the United States. There is again a clear positive correlation between the two, albeit it is less clear-cut than that for France. One explanation for this discrepancy could be that I have data on all types of renters in the US, not only for apartments currently on the market. Taking long-term tenants into consideration implies a larger impact of rent control or subsidized housing, with differences in the extent of such programs and rules between cities. Moreover, the US has more local autonomy regarding taxes and public services than France. More remote places might be attractive for institutional reasons, which can impact gradients and, indirectly, the estimated prices at the city center. However, overall, the mapping between prices estimated using Airbnb properties on the one hand and long-term rental data on the other seems reasonably good.[28]

`dannonce-par-commune-en-2022/` (downloaded on 2023-02-17.)

[24] This constitutes the complete universe of French communes except for 17 communes in Mayotte. For the large cities of Paris, Marseille, and Lyon, the information is available at the level of arrondissements (neighborhoods). Their hedonic regression accounts for surface area, average surface per room, as well as year, trimester, and source of the ad.

[25] I spatially join the centroids of the communes and arrondissements to the city polygons. I drop an additional six communes, which are within the extent of my cities but do not host any Airbnb that was on the market for at least 100 days.

[26] The data cover 219,773 block groups. 80,550 of these block groups (measured at their centroid) are within one of the 70 US cities in my sample. I further restrict the analysis to the 43,636 block groups that host an Airbnb, which meets the minimum criterium of 100 nights available or reserved.

[27] For each block group, I have information about the fraction of apartments that meet certain brackets in the following categories: bedrooms, units in the building, construction year of the building, and year the tenant moved in. Moreover, I control for the fraction of apartments with a kitchen, with plumbing, and for whether the block group borders an ocean or a big lake.

[28] In the case of the US, there are three outliers in Sandy, West Valley City, and Overland Park. These are three of the very few cases in which the global rules, according to which I delimitate cities, lead to suboptimal outcomes. Overland Park might be more accurately described as part of Kansas City, while Sandy and West Valley City should probably form a single city with Salt Lake City.

4 Second-stage regressions

In the second stage, I regress the logs of the prices obtained in the first stage on log city size while controlling for a list of city characteristics. The literature on agglomeration effects typically uses either population or population density to explain wage differentials across cities. Henderson et al. (2021) test more elaborate density measures, but find that they do not offer a real improvement over simply using population density. I mostly follow Combes et al. (2018) in using the log of population size as my primary variable of interest while controlling for log area. This approach can be seen as an unrestricted version of population density, in which the coefficients of population and area are not coerced to be the opposite of each other. Combes et al. (2018) also provide an economic intuition to this approach: Controlling for area is the equivalent of restricting a city from expanding outwards when it is confronted with a higher population size. Correspondingly, they find city size to increase real estate prices more strongly when they control for area, compared to when they do not.

Table 1: Summary statistics, variables of second-stage regressions

Variable	Mean	SD	Q10	Median	Q90	N
Price of representative apartment	54.10	35.98	23.11	40.50	105.49	733
Population in 2015	2,102,975	3,690,911	359,257	906,728	4,219,852	733
Population in 1975	1,018,667	1,854,413	141,651	473,417	2,029,259	733
Area in km2	404	603	83.81	215	846	733
Compactness	0.72	0.11	0.57	0.74	0.85	733
Elevation in m	327	557	13.40	79.68	1,126	733
Difference to 21.11°C	6.74	4.45	1.39	5.79	13.01	733
Located by ocean or big lake	0.35	0.48	0.00	0.00	1.00	733
Capital	0.16	0.36	0.00	0.00	1.00	733
Airbnbs per 1,000 inhabitants	2.77	4.48	0.16	1.04	7.77	733
In 50 cities with most homicides	0.06	0.24	0.00	0.00	0.00	733
Homicides per 100k (Mex)	36.57	33.22	6.10	27.51	86.43	38
Borders USA (Mex)	0.08	0.27	0.00	0.00	0.00	38

Table 1 presents summary statistics of the variables used for these second-stage regressions. Elevation and temperature also come from the Global Human Settlement project.[29] For temperature, I follow Chauvin et al. (2017) in considering the difference to 21.11°C, which they characterize as

[29] They are included in a dataset that describes their urban centers. In some instances, I split these urban centers into more than one city (see Section 2 and Nöbauer, 2023). While I can precisely compute the area and estimated population size for these divided cities, I have the data on temperature and elevation only for the entire urban centers. In the case of split cities, I assign the values of the underlying urban center to all cities that emerge through these splits.

the "middle ground within the [...] range that is often discussed an ideal for human comfort" (p. 27).[30] I lose one observation (Weihai in China) because of a missing value regarding temperature. Capital is an indicator variable for whether a city is the national capital of its respective country, while "[l]ocated by ocean or big lake" is an indicator for whether a city borders a major water body. The first-stage hedonic regression already includes an indicator variable for whether a property is located close to an ocean or big lake ($\geq 80\text{km}^2$). However, besides influencing the price of individual properties, being located at a shore might also impact how cities as a whole are organized and experienced. I also include a measure of the number of Airbnb properties per 1,000 inhabitants to control for the attractiveness of a city to tourists.

Moreover, using my city polygons, I compute a compactness measure based on Angel et al. (2020). It assesses how much the shape of a city resembles a circle on a scale from 0 to 1. Technically, I compute a circle with the same area as the city itself around each city's centroid and then measure the proportion of the circle that intersects with the shape of the city (the "exchange" measure in Angel et al., 2020). The rationale for this is that accessibility is dependent not only on the size of the area in which a given population is distributed, but also on the form that area takes. A circular area makes it easier to provide a high level of accessibility from many locations than a drawn-out or ramified one. The differences in accessibility can, in turn, affect how much people are willing to pay to live in the city center. Figure A4 shows the measure for four exemplary cities corresponding to the highest compactness value, the 75% quantile, the 25% quantile, and the lowest compactness value in the sample.

4.1 Second-stage results

Table 2 shows the main results of my second-stage regressions. All six specifications include country-fixed effects, which implies that I estimate the elasticity of housing costs with respect to city size from within-country variation. The population coefficient is statistically significant at the 1% level in all OLS specifications. Without any controls, I estimate an elasticity of housing costs with respect to city size of 0.139. Once I control for area, this coefficient increases to 0.164. This implies that the association between population size and housing costs is stronger when cities are not allowed to expand outwards. In that case, every additional person must be absorbed by infill (less green space or vacant plots within the city), vertical growth (taller buildings), or reduced living space per person (smaller housing units or more people per housing unit). The difference between the population coefficients under the two settings is 0.025 without additional controls, but it increases to 0.042 once the other controls are introduced.

Column 4 shows my preferred specification. It reports an elasticity of housing costs with respect to city size of 0.161. In other words, if the population size of a city increases by 10%, housing costs rise by 1.61%. This global estimate is somewhat smaller than the estimates of the elasticity of house prices with respect to city size that Combes et al. (2018) report for France. Their estimates

[30] They separately consider temperature differences in January and July, while I only have annual averages.

Table 2: Main specifications

Dependent Variable:	log(Price of representative apartment at city center)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	IV	IV
log(Population)	0.139***	0.164***	0.119***	0.161***	0.118***	0.153*
	(0.015)	(0.057)	(0.013)	(0.050)	(0.016)	(0.084)
log(Area)		-0.030		-0.050		-0.042
		(0.063)		(0.053)		(0.087)
Compactness			-0.045	-0.054	-0.045	-0.053
			(0.118)	(0.112)	(0.118)	(0.113)
Elevation (100m)			-0.007*	-0.007*	-0.007*	-0.007*
			(0.004)	(0.004)	(0.004)	(0.004)
Difference to 21.11°C			0.002	0.002	0.002	0.002
			(0.007)	(0.007)	(0.007)	(0.007)
By ocean / big lake			0.055	0.054	0.055	0.055
			(0.043)	(0.045)	(0.043)	(0.045)
Capital			0.106*	0.105*	0.108*	0.106*
			(0.057)	(0.057)	(0.058)	(0.059)
Airbnbs per 1,000			0.030***	0.029***	0.030***	0.029***
			(0.006)	(0.006)	(0.006)	(0.006)
Cragg-Donald F-Stat					2,959.6	268.0

Note: The table shows regressions of the estimated price of a representative short-term rental property at the city center on city size and control variables. The units of observation are 733 cities. All specifications include country fixed effects. Population in 2015 is instrumented by population in 1975 for the IV specifications. The parentheses show standard errors, which are clustered by country. The levels of significance are * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

range from 0.176 to 0.305, with their preferred estimate being 0.208. At the same time, their estimated area-unrestricted elasticity is 0.109, which is almost precisely what I find. The fact that the difference manifests itself in the area-restricted elasticity could be consistent with French cities being more limited in vertical growth by stricter regulations than cities elsewhere. However, it is important to keep in mind that my sample consists of cities that are, on average, more than 12 times larger than the cities used by Combes et al. (2018). They also estimate an elasticity that is non-linear in population size and find (area-restricted) estimates as large as 0.288 of a city with one million inhabitants and 0.378 for a city as large as Paris. Comparing these estimates with mine suggests that the housing costs in big French cities increase faster in city size than the housing costs in big cities elsewhere.

Columns 5 and 6 mirror the specifications of columns 3 and 4, but introduce an instrumental

variable approach. I instrument log population in 2015 with log population in 1975. The idea is that today's price level (and other recent developments that affect it) might cause people to move into (or away from) a given city and, might therefore, bias the estimation. At the same time, long-past population counts should be unaffected by it. In applying this strategy, I follow a standard approach introduced by Ciccone and Hall (1996) and used amongst others by Combes et al. (2008) and Combes et al. (2018). Unlike these papers, my work deals with a worldwide sample, and, unfortunately, it is impossible to find ancient population counts on that scale. The advantage of the 1975 population data I use, apart from its existence, is that it is provided by the same source, built using the same principles, and covering the same grid as the 2015 population data. However, 40 years are not enough to alleviate concerns about the instrument's validity. As Chauvin et al. (2017), who use population data from 1980 to construct an IV, I argue that columns 5 and 6 should not be interpreted as more than a robustness check. The point estimates are almost unchanged between columns 3 and 5, with the variable of interest still being statistically significant at the 1% level. When I control for area, the coefficient of log population decreases from 0.161 to 0.153 between the OLS and the IV estimation. It is only statistically significant at the 10% level in the IV setting, compared to the 1% level with OLS.

Concerning the control variables, I estimate capital cities to be about 10.6% more expensive than other cities, with the effect being statistically significant at the 10% level. A higher number of Airbnbs is associated with higher prices. This relation is statistically significant at the 1% level. The predicted housing cost difference between a city at the 25% quantile and a city at the 75% quantile of Airbnb properties is 0.079. A statistically significant (at the 10% level) relation exists between elevation and housing costs. However, given that the average elevation is 328 meters, with a median of 80 meters, this effect is quantitatively small. Moreover, I estimate more compact cities to be cheaper and cities at the seaside to be more expensive beyond the properties close to the shore. However, neither of these effects is statistically significant.

Importance of fixed effects

Table A4 explores the explanatory power of the different sets of variables. The first column regresses the log price of the representative apartment at the city center merely on the logs of population and area. Without country-fixed effects, the direction of the effect switches. This behavior is consistent with the fact that lower income countries tend to have denser cities with less living space per person (Jedwab et al., 2021). Specification (2) consists of the control variables only. The point estimates of the controls go in the same direction as in the full specification, but they are larger, which can be explained by the omission of the country-fixed effects. Log population and log area alone (column 1) and the controls alone (column 2) explain an R^2 of around 0.3.

The R^2 increases to 0.48 in column (3), which includes both the logs of population and area and the controls. Adding controls without country-fixed effects still results in a negative correlation

Table 3: Heterogeneity by country

Dependent Variable:	log(Price of representative apartment at city center)						
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	USA	Eurozone	Russia	China	India	Brazil	Mexico
log(Population)	0.23**	0.33***	0.19	0.11	0.43**	0.05	-0.36**
	(0.10)	(0.08)	(0.18)	(0.12)	(0.16)	(0.09)	(0.17)
log(Area)	-0.07	-0.13	0.09	0.02	-0.25	0.05	0.52**
	(0.12)	(0.10)	(0.22)	(0.13)	(0.19)	(0.11)	(0.19)
Compactness	-0.38**	0.33	-0.30	-0.32	-0.22	0.35	-0.74
	(0.18)	(0.32)	(0.43)	(0.32)	(1.21)	(0.23)	(0.50)
Elevation (100m)	-0.03***	-0.03	-0.03	0.00	0.00	0.00	0.00
	(0.01)	(0.03)	(0.03)	(0.01)	(0.03)	(0.01)	(0.01)
Difference to 21.11°C	0.02***	0.08***	0.00	-0.01*	-0.05	-0.01	-0.02
	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)	(0.01)	(0.02)
By ocean / big lake	-0.12**	-0.03	0.04	0.20**	0.19	0.03	0.08
	(0.05)	(0.06)	(0.07)	(0.08)	(0.27)	(0.07)	(0.14)
Airbnbs per 1,000	0.04***	0.02***	0.05**	0.04***	0.74**	0.10***	0.10***
	(0.01)	(0.00)	(0.02)	(0.01)	(0.34)	(0.03)	(0.02)
Country fixed effects	-	Yes	-	-	-	-	-
Observations	70	76	44	112	31	44	38

Note: The table shows regressions of the estimated price of a representative short-term rental property at the city center on city size and control variables. The units of observation are cities. The parentheses show standard errors clustered by country for specification 2 (eurozone) and heteroscedasticiy robust standard errors for all other specifications. The levels of significance are * p < 0.10, ** p < 0.05, *** p < 0.01.

between population count and price.[31] Specification (4) exclusively contains country-fixed effects and shows that they alone generate a R^2 of 0.74. Columns (5) and (6) mirror specifications (2) and (4) of Table 2. Adding the logs of population and area increases the R^2 to 0.78, while additionally adding controls raises it to 0.81.

Geographic heterogeneity

Table 3 repeats the analysis for the six countries with the highest number of cities in the sample, and for the eurozone. Given the low number of observations, these results should be taken with a grain of salt. Nevertheless they show some interesting regularities.

I estimate a statistically significantly positive elasticity for the United States, India, and the euro-

[31] An obvious omitted control variable in this specification is income. I do not include it since all other second-stage regressions either include country-fixed effects or focus on one particular country. Spatially disaggregated within-country data on income per capita is not readily available on a global scale.

zone. In all cases, the estimated elasticity is higher than for the full global sample, with particularly large effects in the latter two. While the elasticity is also somewhat higher in Russia, the estimate is not statistically significant. My results for China and Brazil show positive estimates that are below the world-average and not statistically significant. Mexico is the only country that completely falls out of line. It shows a statistically significant effect that is negative and large. This implies that city size is negatively correlated with housing costs in the Mexican context. I will come back to this below.

As stated above, Combes et al. (2018) find an elasticity of 0.29 for a city with one million inhabitants in the French context. The average eurozone-city in my sample has a population of 1.07 million, which makes that a good comparison. Estimating a model with country-fixed effects for the eurozone, I find a coefficient of 0.33. It is hard to check whether France is representative for the eurozone. If it is, this finding would suggest that the elasticity of housing costs with respect to city size is indeed increasing in city size. Moreover, the estimated elasticity from Combes et al. (2018) might be on the higher end of the global spectrum.

There are at least two plausible explanations for the lower and not statistically significant estimates for Russia, China, and Brazil and for the fact that the coefficient of log area is estimated to be positive. First, these countries might have less stringent regulations concerning building heights or building over green spaces. They might also simply have more room for infill. Second, the results might also be biased towards zero because of data quality issues. While China has the highest number of cities in the sample, delineating the cities and setting their center was harder than anywhere else. For example, the maps displayed by Google Maps are not superimposable to satellite images for Chinese cities because of government regulations. Instead they are shifted in a non-monotonic way (Fuentes, 2019). In contrast, the United States and the eurozone have higher numbers of Airbnbs in the sample than all other countries mentioned in Table 3. This might lead to more precisely estimated housing costs in the first-stage.

While Table 3 does not include capital dummies, Table A5 repeats the analysis without the two largest cities for each entity.[32] Most of the results are very robust to the exclusion of these cities. The notable exception is Russia, where the coefficient drops from 0.19 to 0.02 after the exclusion of Moscow and Saint Petersburg, implying that the positive relation between city size and housing costs is entirely driven by these two metropolises.

Comparison with Chauvin et al. (2017)

Chauvin et al. (2017) also provide recent estimates of the elasticity of housing costs with respect to city size for multiple countries. They focus on other aspects of the spatial equilibrium, amongst others on agglomeration economies. However, an appendix to their paper includes such estimates for the United States, Brazil, China, and India. They estimate their regressions using OLS and IV

[32] The excluded cities are New York, Los Angeles, Paris, Madrid, Moscow, Saint Petersburg, Shanghai, Beijing, Delhi, Mumbai, São Paulo, Rio de Janeiro, Mexico City, and Guadalajara.

specifications, based on data from 2010 and with the population in 1980 as an instrument.[33] This is close enough to the years I use (population data from 2015, short-term rental prices from 2018-19, population from 1975 as an instrument) to expect similar results.

Table A6 presents this comparison. All point estimates correspond to the effect of log population. My preferred OLS specification includes the same controls as Table 3 but excludes area to be consistent with the estimates from Chauvin et al. (2017).[34] My preferred IV specification additionally instruments log population in 2015 with log population in 1975. Our results are very similar for the US, where the data availability is the best.[35] Chauvin et al. (2017) present a specification with log rent and another one with log price as the dependent variable, and my estimates fall right in between the two. I get somewhat lower point estimates for Brazil while confirming the positive and statistically significant elasticity of housing costs with respect to city size. They also report two separate regressions for China. My preferred estimates are again between the two estimates of Chauvin et al. (2017).

The similarity of the results disappears in the case of India. Chauvin et al. (2017) find no statistically significant effect of city size on housing rents for India, with point estimates narrowly above and below zero. They do find agglomeration economies for India that are about 50% higher than for the US, which implies that real wages must increase in city size. They explain this with low migration rates and geographical differences in the level of education but also acknowledge that the data quality of their rent data might offer another explanation (p. 32). In contrast, I do find a statistically significant positive relation between short-term rental prices and city size. The corresponding coefficient is about 35% larger than that for the United States. If amenities increase less (or decrease more) with city size than in the US, this could very well be in line with the standard spatial equilibrium model whose applicability to the Indian context is challenged by Chauvin et al. (2017).

There are some notable methodological differences concerning the estimation of this elasticity. First, Chauvin et al. (2017) do not report to account for property-level characteristics, while my rental price indices are the outcome of hedonic regressions. Second, I estimate prices at the city center, while Chauvin et al. (2017) appear to use city fixed effects without accounting for any geographical within-city dimension. They also do not use city-level controls in the second stage. If I adjust my methodology concerning the second-stage regression, I also estimate elasticities for India that are very small and statistically indistinguishable from zero. While the same pattern emerges for the US and China, albeit to a smaller degree, the results for Brazil are unchanged (strips 3 and 4 of Table

[33] They also include IV estimates based on older population counts: 1900 for the United States, 1920 for Brazil, 1950 for China, and 1951 for India. I restrict my comparison to their first set of IV estimates as it provides better comparability.

[34] They also report coefficients of regressions using density as the independent variable. Those results are qualitatively similar, except for house prices in China, where they report statistically significant effects of around 0.22.

[35] Chauvin et al. (2017) work with household level data and report the following sample sizes: 24.4 mio / 44 mio (rent/price) for the US, 818 k for Brazil, 6.7 k / 25 k (rent/price) for China, and 3.3 k for India. My work builds on the following numbers of Airbnb properties that were available at least 100 out of 365 nights: 282 k for the US, 36 k for Brazil, 169 k for China, and 9.4 k for India.

A6).

Homicides and the elasticity of housing costs with respect to city size

Why are the results that different for Mexico? One possible explanation could be the level of crime. Glaeser and Sacerdote (1999) argue that "it is ironic that the same urban advantages, lower transport costs, faster urban information flows, and the same scale economies that help to make cities more productive also increase the level of crime in the city" (p. 241). Safety considerations might not affect the attractiveness of cities too much when the overall level of crime is low. However, Mexico is in a drug war and experienced over 72,000 homicides in 2018 and 2019 alone.[36] It seems plausible that population density can seem frightening in such an environment. In a Roback (1982) type setting, crime can act as a negative amenity, with crime-ridden places having to offer higher wages or lower real estate prices in equilibrium. If the probability of becoming the victim of a crime increases in city size, this can explain why the positive relationship between city size and real estate prices might not hold in places with high crime rates.

Table 4 explores this dimension. Column (1) reports the baseline regression for the 38 Mexican cities in my sample.[37] I include a dummy for whether a city borders the United States since I expect the Mexican real estate market and potentially also crime rates to be affected by proximity to the US. Column (2) includes the homicide rate per 100,000 inhabitants and an interaction term between the homicide rate and population size. The data originate from the "Instituto Nacional de Estadística y Geografía" and are cleaned and made available by Diego Valle-Jones.[38] I use the average of the 2018 and 2019 homicide rates. The yearly average homicide rate among the 38 Mexican cities in my sample is 36.6 per 100,000 inhabitants (see Table A4), while the average population size of these cities is 1.46 million.

I find a negative interaction between homicides and population size that is statistically significant at the 10% level. This result implies that the negative correlation between short-term rental prices and city size that I find in the Mexican context is particularly strong for cities with a high homicide rate. While controlling for the homicide rate makes the baseline effect of population size smaller and statistically insignificant, its point estimate is still negative. However, if crime is indeed a driver of this reverse effect, it is plausible that even the safer cities in Mexico are affected to some degree.

I, therefore, try to go beyond Mexico. Column 3 of Table 4 replicates the baseline specification for the entire worldwide sample. Column 4 adds an indicator for whether a city appears in the 2018 or 2019 versions of the list of the 50 cities with the highest homicide rates that is published each year by the "Consejo Ciudadano para la Seguridad Pública y la Justicia Penal AC".[39] The baseline

[36] This number is based on the same data from the INEGI that I describe below.

[37] The crime data are based on metro areas and includes several big municipios that are not part of a metro area. For details, see `https://github.com/diegovalle/mxmortalitydb`, last accessed on 2023-06-11. I match these metro areas to my cities by name and verify that they include the same center. However, the boundaries of the metro areas and my cities differ to some extent.

[38] I last downloaded the data on 2023-06-09 from `https://github.com/diegovalle/mxmortalitydb`.

[39] I downloaded the rankings on 2023-05-27 from `https://geoenlace.net/seguridadjusticiaypaz/webpage/`

Table 4: Interaction with the homicide rate

Dependent Variable:	log(Price of representative apartment at city center)			
Extent:	Mexico	Mexico	World	World
Model:	(1)	(2)	(3)	(4)
log(Population)	-0.284	-0.201	0.161***	0.164***
	(0.191)	(0.189)	(0.050)	(0.050)
log(Area)	0.419*	0.451**	-0.050	-0.050
	(0.215)	(0.218)	(0.053)	(0.054)
Compactness	-0.820	-0.919	-0.054	-0.050
	(0.491)	(0.581)	(0.112)	(0.113)
Elevation (100m)	0.002	0.004	-0.007*	-0.007*
	(0.006)	(0.006)	(0.004)	(0.004)
Difference to 21.11°C	-0.007	-0.012	0.002	0.002
	(0.020)	(0.019)	(0.007)	(0.007)
By ocean / big lake	0.058	0.089	0.054	0.054
	(0.128)	(0.145)	(0.045)	(0.044)
Capital			0.105*	0.100*
			(0.057)	(0.058)
Airbnbs per 1,000	0.106***	0.106***	0.029***	0.029***
	(0.022)	(0.024)	(0.006)	(0.006)
Borders USA	0.254*	0.382***		
	(0.149)	(0.118)		
Homicides per 100k		0.051*		
		(0.025)		
log(Population) × Homicides per 100k		-0.004*		
		(0.002)		
In 50 most homicides				0.961***
				(0.339)
log(Population) × In 50 most homicides				-0.069***
				(0.024)
Country fixed effects	-	-	Yes	Yes
Observations	38	38	733	733

Note: The table shows regressions of the estimated price of a representative short-term rental property at the city center on city size and control variables. The units of observation are cities. The parentheses show heteroscedasticity robust standard errors in models (1) and (2) and standard errors clustered by country in models (3) and (4). The levels of significance are * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

effect of being in this ranking is positive, which reflects the fact that most of these cities are located in upper-middle-income countries.[40] However, there is a negative interaction between city size and being one of the cities with the highest homicide rates. The corresponding regression coefficient is statistically significant at the 1% level. I estimate the elasticity of housing costs with respect to city size to decrease from 0.164 to 0.095 for the most dangerous cities according to this definition.

5 Conclusion

In this paper, I estimate the elasticity of housing costs with respect to city size. I conduct the analysis on a worldwide scale, using 733 cities with at least 300,000 inhabitants and 100 Airbnb properties. I am able to work on this international scale because I use novel data on short-term rental properties from Airbnb as a proxy for housing costs. In a first-stage hedonic regression, I estimate the price of a representative property at the center of each city. I then use these prices in a second-stage regression, regressing them on population, area, city-level controls, and country-fixed effects.

My preferred estimate of the elasticity of housing costs with respect to city size is 0.16. This is somewhat less than the estimate of 0.21 that Combes et al. (2018) find for a sample of French cities that are more than 12 times smaller on average. When not controlling for area, both our samples yield an estimate of 0.11. I find the elasticity to differ substantially by country/region, estimating a coefficient of 0.33 for the eurozone. This is in line with non-linear estimates that Combes et al. (2018) provide for a hypothetical French city with one million inhabitants (0.29) and for Paris (0.38). Assuming that French cities and other cities in the eurozone are alike, this supports their finding of a non-linear elasticity of housing costs with respect to city size.

It also suggests that large eurozone cities face above-average elasticities of housing costs with respect to city size compared to other large cities worldwide, as do cities in India. In particular, I estimate elasticities that are positive but considerably smaller and not statistically significant for Russia (in particular without Moscow and Saint Petersburg), China, and Brazil. Especially given the small sample sizes of these country regressions, I cannot rule out that data issues drive part of this discrepancy. However, given that I control for city area, I hypothesize that stricter building height regulations in the eurozone and India might play some role, by limiting how much the housing stock can adjust as a reaction to population growth. Infill development within the existing boundaries of a city can play a similar role of adjustment, where cities in the eurozone and India might contain fewer empty plots.

An alternative and perhaps complementary explanation is based on the Rosen-Roback model (Rosen, 1979; Roback, 1982). The model predicts that differences in wages, housing costs, and

archivos.php.

[40] To classify countries by income, I use the World Bank definitions, available at `https://datahelpdesk.worldbank.org/knowledgebase/articles/906519-world-bank-country-and-lending-groups` (last accessed: 2023-06-09).

urban amenities counterbalance each other. If a city has high wages and low housing costs, it will most likely not have good amenities. Otherwise, residents from other cities would move, driving housing costs up and wages down. Each of the three factors has an elasticity with respect to city size. Concerning the elasticity of wages with respect to city size, empirical research tends to find larger agglomeration effects in developing countries than in high-income countries (Chauvin et al., 2017; Henderson et al., 2021). Combining this evidence with my findings on urban costs has implications about the elasticity of amenities with respect to city size. If wages in the eurozone increase less in city size and housing costs increase more, the Rosen-Roback model would predict that the quality of amenities increases more strongly in city size in the eurozone (or decreases less strongly in city size) than elsewhere.

Finally, the suggestive evidence of a negative elasticity of housing costs with respect to city size for Mexico, and the negative interaction between that elasticity and homicides, both in Mexico and worldwide, open room for future research. While Glaeser and Sacerdote (1999) find crime to increase in city size for the United States, Ahlfeldt and Pietrostefani (2019) find it to decrease in density in other OECD countries. Which types of crime affect other urban costs (and perhaps also benefits) in which contexts remains an exciting open question.

6 Bibliography

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

Akbar, P.A., Couture, V., Duranton, G., and Storeygard, A., 2021. Mobility and congestion in urban India, Working Paper.

Alonso, W., 1964. *Location and Land Use. Toward a General Theory of Land Rent*, Harvard University Press.

Angel, S., Franco, S.A., Liu, Y., and Blei, A.M., 2020. The shape compactness of urban footprints, *Progress in Planning*, 139, 100429.

Chauvin, J.P., Glaeser, E., Ma, Y., and Tobio, K., 2017. What is different about urbanization in rich and poor countries? Cities in Brazil, China, India and the United States, *Journal of Urban Economics*, 98, 17–49.

Ciccone, A. and Hall, R., 1996. Productivity and the density of economic activity, *American Economic Review*, 86 (1), 54–70.

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

Combes, P.P., Duranton, G., and Gobillon, L., 2018. The costs of agglomeration: House and land prices in French cities, *The Review of Economic Studies*, 86 (4), 1556–1589.

Combes, P.P. and Gobillon, L., 2015. The empirics of agglomeration economies, *in:* G. Duranton, J.V. Henderson, and W.C. Strange, eds., *Handbook of Regional and Urban Economics*, Elsevier, *Handbook of Regional and Urban Economics*, vol. 5, 247–348.

Duranton, G. and Puga, D., 2004. Micro-foundations of urban agglomeration economies, *in:* J.V. Henderson and J.F. Thisse, eds., *Cities and Geography*, Elsevier, *Handbook of Regional and Urban Economics*, vol. 4, 2063–2117.

Duranton, G. and Puga, D., 2015. Urban land use, *in:* G. Duranton, J.V. Henderson, and W.C. Strange, eds., *Handbook of Regional and Urban Economics 5*, Elsevier, 467–560.

FIFA, 2010. 2018 FIFA World Cup™ bid evaluation report: Russia, Report, Fédération Internationale de Football Association.

Florczyk, A., Corbane, C., Schiavina, M., Pesaresi, M., Maffenini, L., Melchiorri, M., Politis, P., Sabo, F., Freire, S., Ehrlich, D., Kemper, T., Tommasi, P., Airaghi, D., and Zanchetta, L., 2019. GHS Urban Centre Database 2015, multitemporal and multidimensional attributes, R2019A, European Commission, Joint Research Centre (JRC). Dataset.

Fuentes, E., 2019. Why GPS coordinates look wrong on maps of China, `https://www.serviceobjects.com/blog/why-gps-coordinates-look-wrong-on-maps-of-china/` (last accessed: 13.01.2023).

Gallagher, L., 2017. *The Airbnb Story: How Three Ordinary Guys Disrupted an Industry, Made Billions . . . and Created Plenty of Controversy*, Harper Business.

Glaeser, E., 2011. *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*, The Penguin Press.

Glaeser, E.L. and Sacerdote, B., 1999. Why is there more crime in cities?, *Journal of Political Economy*, 107 (6), 225–258.

Henderson, J.V., 1974. The sizes and types of cities, *The American Economic Review*, 64 (4), 640–656.

Henderson, J.V., Nigmatulina, D., and Kriticos, S., 2021. Measuring urban economic density, *Journal of Urban Economics*, 125, 103188.

Henderson, V., 2002. Urban primacy, external costs, and quality of life, *Resource and Energy Economics*, 24 (1), 95–106.

Jedwab, R., Barr, J., and Brueckner, J.K., 2022. Cities without skylines: Worldwide building-height gaps and their possible determinants and implications, *Journal of Urban Economics*, 132, 103507.

Jedwab, R., Loungani, P., and Yezer, A., 2021. Comparing cities in developed and developing countries: Population, land area, building height and crowding, *Regional Science and Urban Economics*, 86, 103609.

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

Muth, R.F., 1969. *Cities and Housing*, University of Chicago Press.

Nöbauer, B., 2023. The monocentric city model worldwide: Rent, density, and transport cost gradients in 734 cities, Working Paper.

Pope, D.G. and Pope, J.C., 2012. Crime and property values: Evidence from the 1990s crime drop, *Regional Science and Urban Economics*, 42, 177–188.

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

Rosen, S., 1979. Wage-based indexes of urban quality of life, *in:* P. Mieszkowski and M. Straszheim, eds., *Current Issues in Urban Economics*, John Hopkins University Press, 74–104.

Schiavina, M., Freire, S., and MacManus, K., 2019. GHS population grid multitemporal (1975, 1990, 2000, 2015) R2019A, European Commission, Joint Research Centre (JRC). Dataset.

Shirk, D. and Wallman, J., 2015. Understanding Mexico's drug violence, *Journal of Conflict Resolution*, 59 (8), 1348–1376.

von Briel, D. and Dolnicar, S., 2021. The evolution of Airbnb regulations, *in:* S. Dolnicar, ed., *Airbnb before, during and after COVID-19*, University of Queensland.

Figure A1: Control variables of hedonic regression

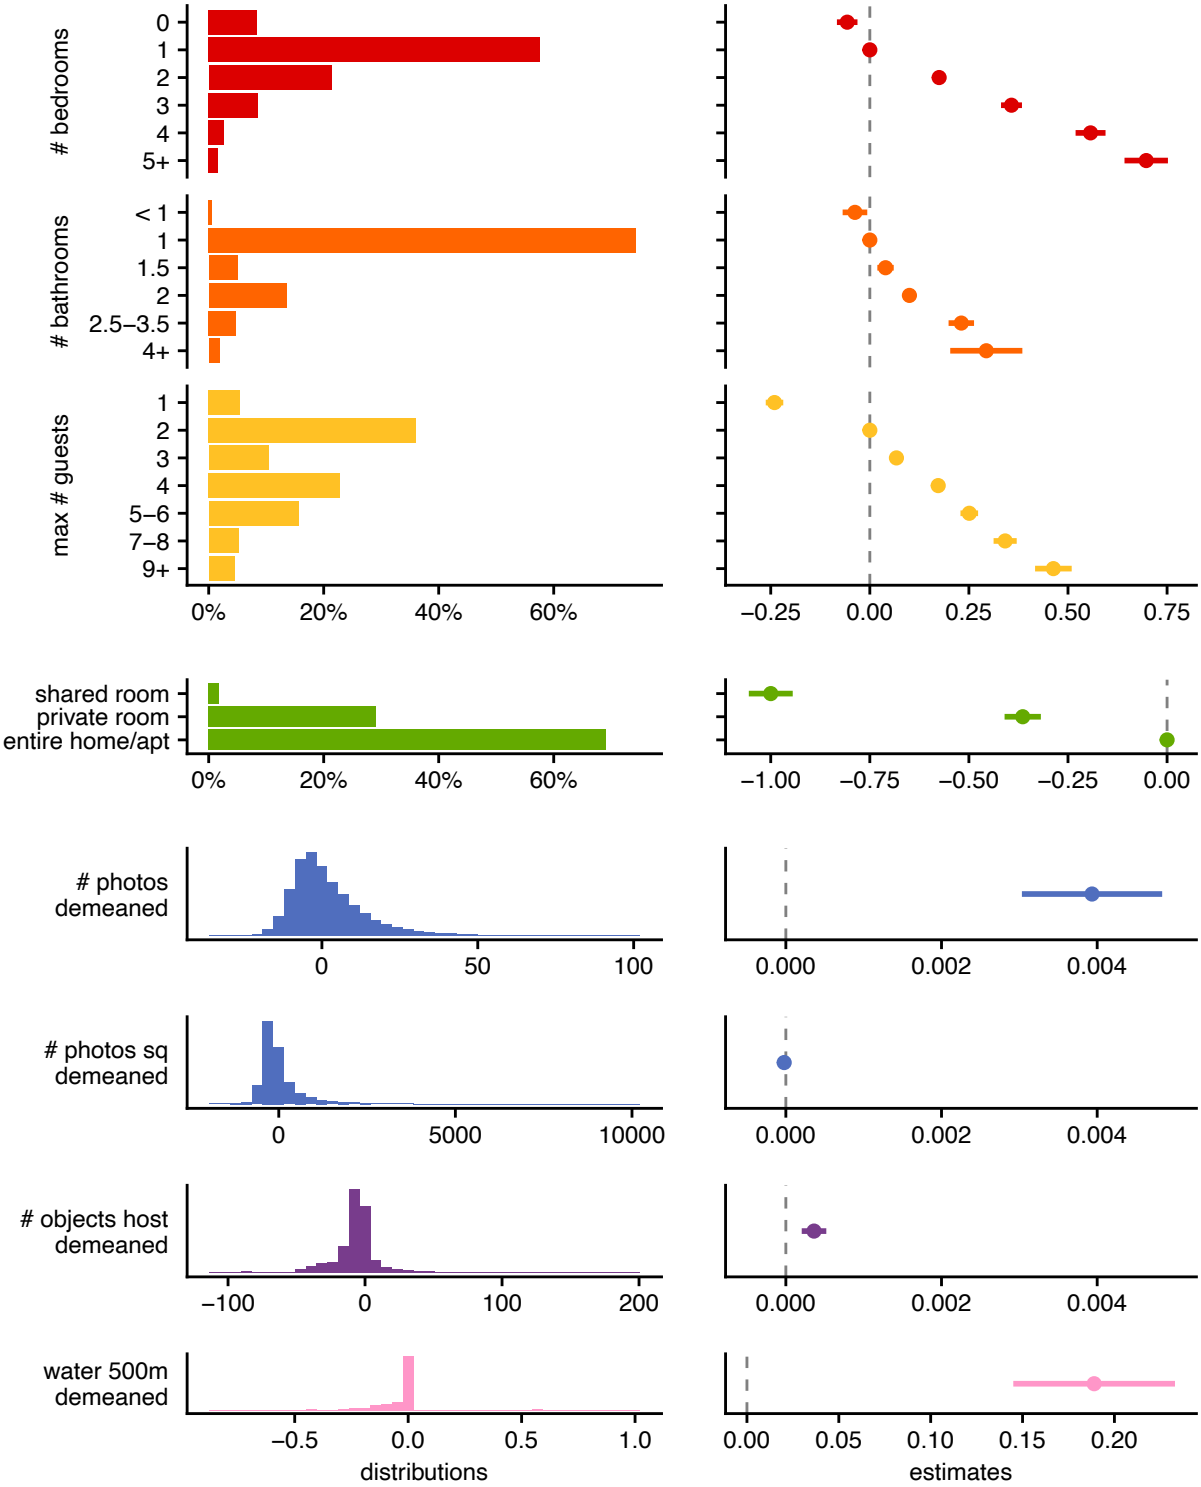

Note: The left column of the figure shows the distributions of all hedonic first-stage regression control variables not classified as "amenities". For the last four panels of this column, the variables are demeaned by subtracting their respective city average. The right column shows the corresponding regression coefficients, with 99% confidence intervals based on standard errors that are clustered by city. The number of observations used in the regression is 1,532,862.

Figure A2: Control variables of hedonic regression: amenities

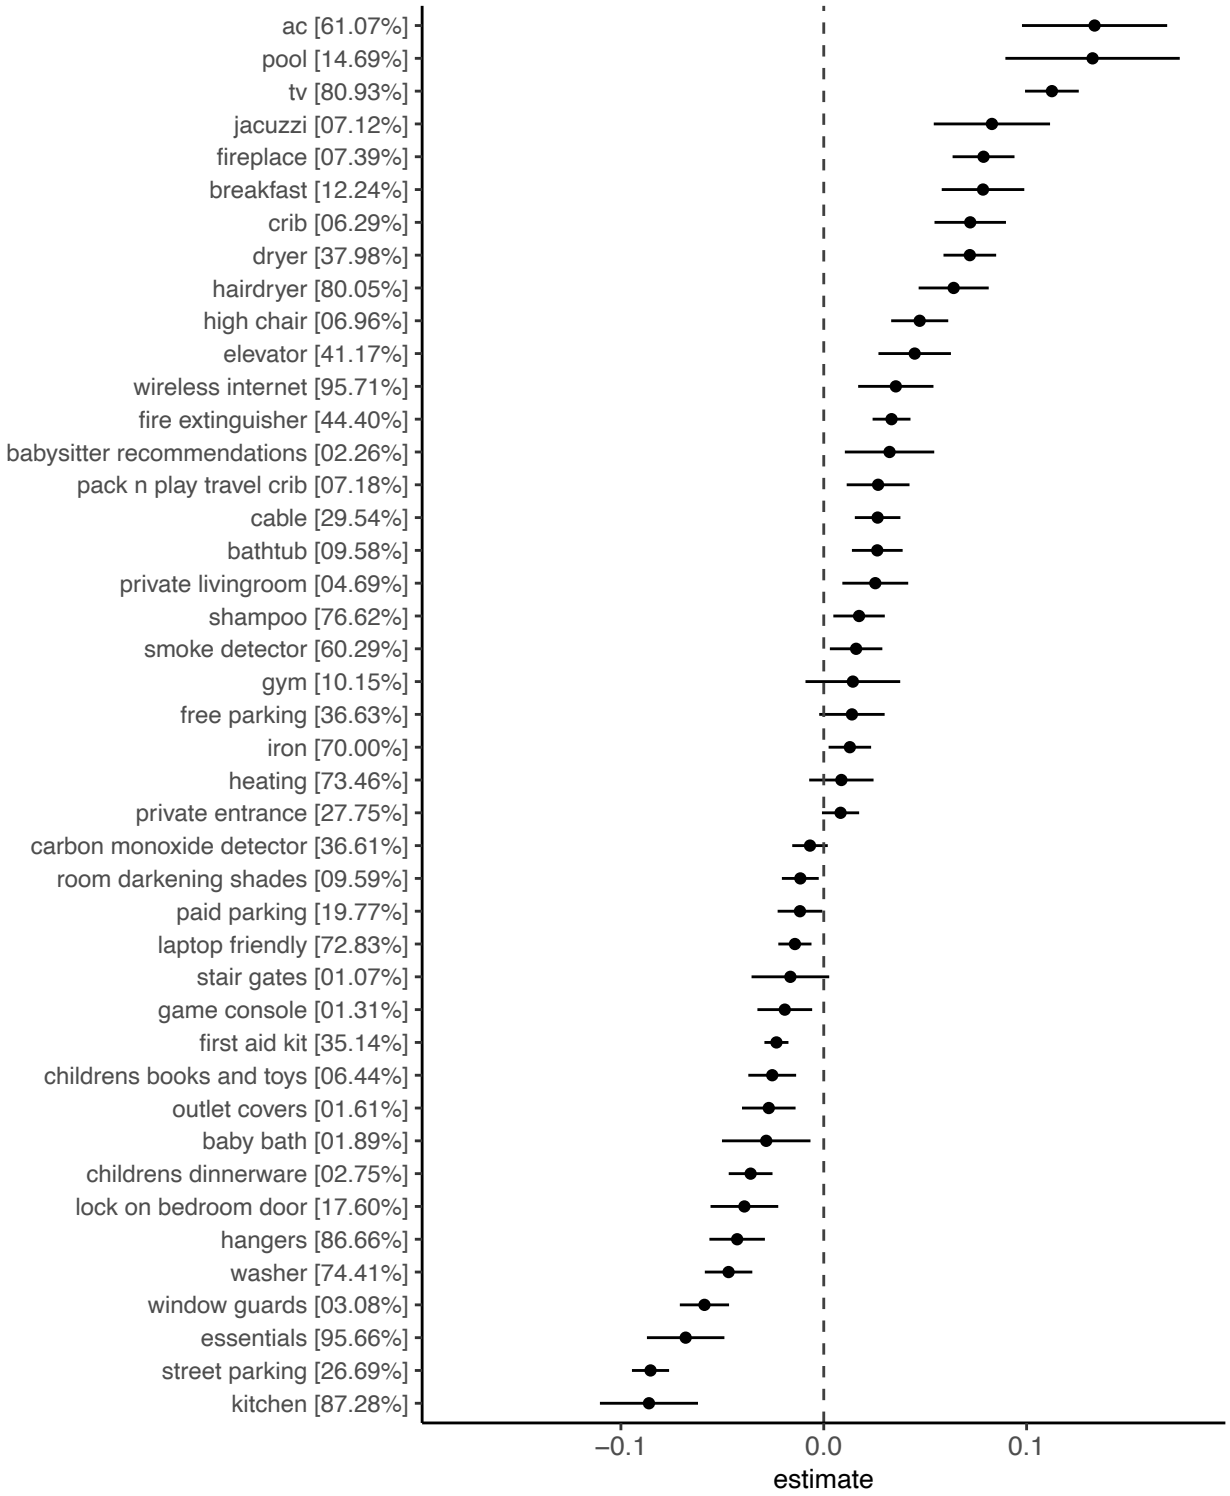

Note: The y-axis shows all 43 amenities which are available in at least 1% of the Airbnb properties in the sample, with the proportion of properties in which the respective amenity is available in brackets. I demean all of these amenities by city and include them as control variables in the first-stage hedonic regression. The figure shows the estimated coefficients of these variables, together with 99% confidence intervals based on standard errors that are clustered by city. The number of observations used in the regression is 1,532,862.

Table A1: Estimated nightly short-term rental rate of a representative apartment at the city center in USD

1	Amsterdam	NLD	252.26
2	San Francisco	USA	242.82
3	London	GBR	230.98
4	New York	USA	224.78
5	Austin	USA	186.78
6	Boston	USA	182.45
7	Seattle	USA	166.25
8	Washington	USA	163.31
9	Copenhagen	DNK	161.39
10	Miami	USA	160.57
11	Paris	FRA	158.82
12	Dubai	ARE	158.47
13	Kuwait City	KWT	157.47
14	Portland	USA	157.11
15	Stockholm	SWE	151.17
16	Dublin	IRL	150.95
17	Chicago	USA	149.43
18	Zürich	CHE	148.15
19	Sydney	AUS	146.44
20	Las Vegas	USA	140.62
21	Long Branch	USA	140.52
22	Roma	ITA	138.42
23	München	DEU	136.79
24	Edinburgh	GBR	136.58
25	New Orleans	USA	135.81
26	Oakland	USA	134.11
27	San Diego	USA	131.99
28	Denver	USA	131.71
29	Honolulu	USA	131.57
30	Lagos	NGA	131.31
31	Vancouver	CAN	131.17
32	Milano	ITA	130.26
33	Göteborg	SWE	130.21
34	Atlanta	USA	127.99
35	Columbus	USA	127.88
36	Oslo	NOR	125.63
37	San Juan	PRI	125.18
38	Cardiff	GBR	125.07
39	Sarasota	USA	123.59
40	Frankfurt am Main	DEU	123.48
41	Genève	CHE	123.45
42	Brighton	GBR	123.33
43	Detroit	USA	122.85
44	Liverpool	GBR	122.40
45	Basel	CHE	121.79
46	Cleveland	USA	121.79
47	Tokyo	JPN	121.41
48	Grand Rapids	USA	120.92
49	Toronto	CAN	120.58
50	Berlin	DEU	118.12
51	Indianapolis	USA	118.09
52	Louisville	USA	115.67
53	Denpasar	IDN	115.49
54	Singapore	SGP	115.28
55	San Antonio	USA	115.12
56	Charlotte	USA	115.11
57	Philadelphia	USA	114.56
58	Los Angeles	USA	114.50
59	Hong Kong	CHN	113.72
60	Pittsburgh	USA	113.19
61	Hamburg	DEU	111.48
62	Houston	USA	111.40
63	Kinshasa	COD	110.87
64	Utrecht	NLD	110.78
65	Firenze	ITA	109.86
66	Manchester	GBR	109.49
67	Barcelona	ESP	108.27
68	Kyoto	JPN	108.14
69	Fort Worth	USA	107.80
70	Jerusalem	ISR	106.49
71	Sacramento	USA	106.39
72	Milwaukee	USA	105.76
73	Palma	ESP	105.69
74	Québec	CAN	105.55
75	Tel Aviv	ISR	105.27
76	Hialeah	USA	104.85
77	Fort Lauderdale	USA	104.58
78	Cincinnati	USA	103.13
79	Minneapolis	USA	102.84
80	Bristol	GBR	102.39
81	Dallas	USA	100.34
82	Wien	AUT	100.23
83	Providence	USA	99.98
84	Kansas City	USA	99.60
85	Tucson	USA	99.30
86	Phoenix	USA	98.95
87	Orlando	USA	98.44
88	Baltimore	USA	98.31
89	Oklahoma City	USA	98.30
90	Leiden	NLD	97.63
91	Auckland	NZL	97.52
92	San Jose	USA	97.50
93	Southampton	GBR	97.41
94	Köln	DEU	97.37
95	Omaha	USA	97.22

96	Concord	USA	96.71
97	Belfast	GBR	96.14
98	Rotterdam	NLD	95.59
99	Den Haag	NLD	95.52
100	Manama	BHR	95.39
101	Jacksonville	USA	95.38
102	Des Moines	USA	95.35
103	Helsinki	FIN	95.23
104	Tampa	USA	95.07
105	Montréal	CAN	94.38
106	Bologna	ITA	94.23
107	Saint Louis	USA	94.10
108	Beijing	CHN	93.82
109	Bordeaux	FRA	93.40
110	Leeds	GBR	93.23
111	Bilbao	ESP	92.10
112	Sevilla	ESP	91.46
113	Praha	CZE	90.99
114	Bruxelles	BEL	90.94
115	Adelaide	AUS	90.38
116	Memphis	USA	90.04
117	Fresno	USA	89.50
118	Düsseldorf	DEU	89.38
119	Colorado Springs	USA	88.78
120	Melbourne	AUS	88.70
121	Glasgow	GBR	88.66
122	Antwerpen	BEL	88.55
123	Saint Petersburg	USA	88.48
124	Salt Lake City	USA	88.17
125	Nürnberg	DEU	87.71
126	Stuttgart	DEU	86.97
127	Madrid	ESP	86.65
128	Strasbourg	FRA	86.27
129	Newcastle upon Tyne	GBR	86.23
130	Birmingham	GBR	86.02
131	Moscow	RUS	85.98
132	Port of Spain	TTO	85.33
133	Dayton	USA	84.90
134	Cartagena	COL	84.87
135	Osaka	JPN	84.67
136	Lyon	FRA	83.63
137	Kolkata	IND	83.34
138	Leicester	GBR	83.20
139	Sapporo	JPN	83.03
140	Rochester	USA	82.18
141	Bakersfield	USA	82.06
142	Ottawa	CAN	81.71
143	Norfolk	USA	81.69
144	Nottingham	GBR	81.48
145	Cape Town	ZAF	81.01
146	Jeonju	KOR	80.89
147	Gold Coast	AUS	80.84
148	Nice	FRA	80.81
149	Accra	GHA	80.38
150	Buffalo	USA	80.02
151	Hannover	DEU	79.29
152	Luanda	AGO	78.85
153	Brisbane	AUS	78.22
154	Sheffield	GBR	77.85
155	Lille	FRA	77.51
156	Kingston	JAM	77.27
157	Aurora	USA	77.07
158	Bergamo	ITA	77.05
159	Xiamen City	CHN	77.01
160	Tulsa	USA	76.91
161	Tainan	TWN	76.72
162	Lisboa	PRT	76.09
163	Sendai	JPN	75.77
164	Ogden	USA	75.64
165	Málaga	ESP	75.59
166	Kanazawa	JPN	75.26
167	Augsburg	DEU	74.98
168	Calgary	CAN	74.95
169	Albuquerque	USA	74.83
170	Asahikawa	JPN	74.43
171	Bloemfontein	ZAF	74.24
172	Hull	GBR	74.04
173	Bonn	DEU	73.72
174	Tallinn	EST	73.64
175	Fukuoka	JPN	72.94
176	Marrakesh	MAR	72.63
177	Portsmouth	GBR	71.51
178	Bremen	DEU	71.40
179	Dresden	DEU	70.67
180	Overland Park	USA	70.67
181	Leipzig	DEU	70.11
182	Naha	JPN	69.65
183	Porto	PRT	69.21
184	Zhoushan	CHN	69.08
185	Winnipeg	CAN	68.93
186	Perth	AUS	68.73
187	Torino	ITA	68.65
188	Coventry	GBR	68.21
189	Ensenada	MEX	68.08
190	Beirut	LBN	67.72
191	Bangkok	THA	67.36
192	Seoul	KOR	67.34
193	Stockton	USA	67.31

194	Haiphong	VNM	67.10	243	Marseille	FRA	58.73
195	Hangzhou	CHN	66.98	244	Athens	GRC	58.58
196	Granada	ESP	66.17	245	Bayuquan	CHN	58.48
197	Napoli	ITA	65.97	246	Qingdao	CHN	58.42
198	Chaozhou	CHN	65.89	247	Abu Dhabi	ARE	57.72
199	Suzhou	CHN	65.88	248	Shenzhen	CHN	57.70
200	London	CAN	65.64	249	Kraków	POL	57.33
201	Liège	BEL	65.59	250	Panamá	PAN	57.32
202	Dortmund	DEU	65.46	251	Windhoek	NAM	57.15
203	Toulouse	FRA	65.22	252	Vilnius	LTU	56.49
204	València	ESP	64.46	253	Dali	CHN	56.30
205	Edmonton	CAN	64.36	254	Kobe	JPN	56.23
206	Taipei	TWN	64.33	255	Viña del Mar	CHL	55.85
207	Sjanghai	CHN	64.23	256	Surrey	CAN	55.04
208	Genova	ITA	64.21	257	Hiroshima	JPN	55.01
209	Mannheim	DEU	63.93	258	Zagreb	HRV	54.87
210	Lahore	PAK	63.48	259	Kaohsiung	TWN	54.85
211	Fez	MAR	63.36	260	Bratislava	SVK	54.77
212	Datong	CHN	63.30	261	Gelsenkirchen	DEU	54.35
213	Budapest	HUN	63.29	262	Ahmedabad	IND	54.11
214	Haifa	ISR	63.22	263	Istanbul	TUR	53.78
215	Nantes	FRA	63.21	264	Brno	CZE	53.69
216	Matsuyama	JPN	63.05	265	Tijuana	MEX	53.45
217	Doha	QAT	62.93	266	Shillong	IND	53.42
218	Zaragoza	ESP	62.45	267	Casablanca	MAR	53.18
219	Kitakyushu	JPN	62.31	268	Alicante	ESP	52.96
220	Rouen	FRA	62.24	269	Grenoble	FRA	52.93
221	Wenzhou	CHN	61.84	270	Toulon	FRA	52.83
222	Sanya	CHN	61.79	271	Zhaoqing	CHN	52.51
223	Mumbai	IND	61.70	272	Kampala	UGA	52.36
224	Agadir	MAR	61.69	273	Mérida	MEX	52.05
225	Stoke-on-Trent	GBR	61.60	274	Sharjah	ARE	52.02
226	Kumamoto	JPN	61.43	275	Bydgoszcz	POL	51.88
227	Abidjan	CIV	61.33	276	Colombo	LKA	51.80
228	Takamatsu	JPN	61.18	277	New Taipei	TWN	51.60
229	Katowice	POL	61.11	278	Taichung	TWN	51.59
230	Bochum	DEU	60.88	279	Gdansk	POL	51.33
231	Huancayo	PER	60.87	280	Wroclaw	POL	51.13
232	Essen	DEU	60.84	281	Douala	CMR	50.93
233	Kitchener	CAN	60.71	282	San Pedro Sula	HND	50.65
234	Bari	ITA	60.68	283	Kyiv	UKR	50.56
235	Petah Tikva	ISR	60.61	284	Murcia	ESP	50.51
236	Nagoya	JPN	60.54	285	Yangzhou	CHN	50.43
237	Saint Petersburg	RUS	59.93	286	Addis Ababa	ETH	50.42
238	Riga	LVA	59.83	287	Durban	ZAF	50.39
239	Shaoxing	CHN	59.66	288	Xiangyang	CHN	50.30
240	Las Palmas	ESP	59.37	289	Amman	JOR	49.90
241	Nairobi	KEN	59.26	290	Dalian	CHN	49.81
242	Vladivostok	RUS	59.12	291	Warszawa	POL	49.75

292	Valledupar	COL	49.54	341	Ningbo	CHN	43.38
293	Duisburg	DEU	49.32	342	Buenos Aires	ARG	43.29
294	Hsinchu	TWN	48.97	343	Minsk	BLR	43.13
295	Poznan	POL	48.68	344	Rabat	MAR	42.96
296	Guangzhou	CHN	48.65	345	Port Elizabeth	ZAF	42.87
297	Lublin	POL	48.44	346	Kuala Lumpur	MYS	42.75
298	Oaxaca	MEX	48.13	347	Kazan	RUS	42.49
299	Palermo	ITA	47.98	348	Changzhou	CHN	42.48
300	Catania	ITA	47.92	349	Plovdiv	BGR	42.45
301	Lusaka	ZMB	47.80	350	Rio de Janeiro	BRA	42.40
302	Santa Cruz d. Tenerife	ESP	47.70	351	Harare	ZWE	42.37
303	Prayagraj	IND	47.57	352	Kagoshima	JPN	42.34
304	Mazatlán	MEX	47.11	353	Nha Trang	VNM	42.25
305	Johannesburg	ZAF	47.05	354	Antalya	TUR	42.15
306	Kigali	RWA	46.71	355	Huangdao District	CHN	42.11
307	Taoyuan	TWN	46.66	356	Yaoundé	CMR	42.11
308	Muscat	OMN	46.60	357	San José	CRI	41.90
309	Tianjin	CHN	46.59	358	Sandy	USA	41.82
310	Mombasa	KEN	46.56	359	Dhaka	BGD	41.62
311	Chiayi	TWN	46.47	360	Cuernavaca	MEX	41.50
312	Wuppertal	DEU	46.42	361	Algiers	DZA	41.33
313	Karachi	PAK	46.33	362	Weifang	CHN	41.19
314	Toshkent	UZB	45.94	363	Santo Domingo	DOM	40.73
315	Tangier	MAR	45.87	364	Lodz	POL	40.72
316	Liangshan	CHN	45.87	365	West Valley City	USA	40.59
317	Changsha	CHN	45.76	366	Chiang Mai	THA	40.58
318	Samara	RUS	45.63	367	Delhi	IND	40.50
319	Abuja	NGA	45.59	368	Xining	CHN	40.34
320	Quanzhou	CHN	45.40	369	Florianópolis	BRA	40.13
321	George Town	MYS	45.32	370	Monterrey	MEX	40.13
322	Cotonou	BEN	45.30	371	Tétouan	MAR	39.99
323	Meknes	MAR	45.29	372	Port-au-Prince	HTI	39.85
324	Binhai New Area	CHN	45.13	373	Meilan District	CHN	39.85
325	Huizhou	CHN	45.08	374	Maputo	MOZ	39.83
326	Nizhny Novgorod	RUS	45.05	375	Santiago	DOM	39.70
327	Chengdu	CHN	45.00	376	Nanchang	CHN	39.66
328	Ciudad de México	MEX	44.85	377	Wuhan	CHN	39.62
329	Udaipur	IND	44.76	378	Tangshan	CHN	39.58
330	Bandaraya Melaka	MYS	44.71	379	Ho Chi Minh City	VNM	39.58
331	Huaiyin	CHN	44.70	380	Jaipur	IND	39.57
332	Daegu	KOR	44.49	381	Guayaquil	ECU	39.50
333	Zhanjiang	CHN	44.46	382	Mexicali	MEX	39.39
334	Acapulco	MEX	44.40	383	Santa Marta	COL	39.37
335	Nanjing	CHN	43.98	384	Bucharest	ROU	39.34
336	Omsk	RUS	43.97	385	Zhenjiang	CHN	39.14
337	Thessaloniki	GRC	43.91	386	São Paulo	BRA	39.10
338	Fuzhou	CHN	43.86	387	Concepción	CHL	39.03
339	Chongqing	CHN	43.86	388	Ürümqi	CHN	38.82
340	Temuco	CHL	43.62	389	Cancún	MEX	38.78

390	Busan	KOR	38.69	439	Rostov-on-Don	RUS	35.07
391	Yekaterinburg	RUS	38.65	440	Santa Fe	ARG	34.88
392	Weihai	CHN	38.56	441	Mendoza	ARG	34.84
393	Meizhou	CHN	38.48	442	Jiaxing	CHN	34.82
394	Morelia	MEX	38.42	443	Kenitra	MAR	34.74
395	Baguio	PHL	38.38	444	Sarajevo	BIH	34.63
396	Montevideo	URY	38.27	445	Corrientes	ARG	34.63
397	Yerevan	ARM	38.05	446	Yantai	CHN	34.62
398	Nanning	CHN	38.05	447	Tolyatti	RUS	34.61
399	Wuxi	CHN	38.04	448	Coimbatore	IND	34.60
400	Shijiazhuang	CHN	37.60	449	Rawalpindi	PAK	34.53
401	La Habana	CUB	37.55	450	Lviv	UKR	34.48
402	Gwangju	KOR	37.53	451	Orizaba	MEX	34.20
403	Freetown	SLE	37.47	452	Qinhuangdao	CHN	34.18
404	Santa Cruz d. l. Sierra	BOL	37.40	453	Tver	RUS	34.09
405	Phnom Penh	KHM	37.38	454	Santiago	CHL	34.08
406	Guilin	CHN	37.27	455	Oran	DZA	34.00
407	Lanzhou	CHN	37.27	456	Salta	ARG	33.88
408	Nur-Sultan	KAZ	37.23	457	Makassar	IDN	33.82
409	Antananarivo	MDG	37.14	458	Puebla	MEX	33.76
410	Tampico	MEX	37.13	459	Nangang	CHN	33.70
411	Pukou	CHN	36.99	460	Subang Jaya	MYS	33.68
412	Kandy	LKA	36.80	461	Baku	AZE	33.68
413	Xi'an	CHN	36.72	462	Pohang-si	KOR	33.47
414	Querétaro	MEX	36.69	463	Bengaluru	IND	33.37
415	Daejeon	KOR	36.63	464	Ipoh	MYS	33.37
416	Zhongshan	CHN	36.63	465	Krasnodar	RUS	33.31
417	Cairo	EGY	36.62	466	Izhevsk	RUS	33.31
418	Belgrade	SRB	36.58	467	Cúcuta	COL	33.27
419	Dakar	SEN	36.53	468	Pucallpa	PER	32.86
420	Santos	BRA	36.45	469	Chihuahua	MEX	32.84
421	Rizhao	CHN	36.42	470	Piura	PER	32.84
422	Odesa	UKR	36.16	471	Salvador	BRA	32.67
423	Quanshan	CHN	36.14	472	Tbilisi	GEO	32.66
424	Pune	IND	35.97	473	Liuzhou	CHN	32.56
425	Kota Kinabalu	MYS	35.92	474	Novokuznetsk	RUS	32.55
426	Zanzibar City	TZA	35.88	475	Ciudad Obregón	MEX	32.47
427	Hrodna	BLR	35.82	476	Ulsan	KOR	32.44
428	Kochi	IND	35.80	477	Volgograd	RUS	32.31
429	Maceió	BRA	35.77	478	Celaya	MEX	32.28
430	Sofia	BGR	35.75	479	Darjeeling	IND	32.28
431	Zhuhai	CHN	35.72	480	Barranquilla	COL	32.23
432	Chisinau	MDA	35.45	481	Jingdezhen	CHN	32.20
433	Yichang	CHN	35.44	482	Luoyang	CHN	32.17
434	Praia Grande	BRA	35.44	483	Tula	RUS	32.16
435	Varanasi	IND	35.36	484	Dnipro	UKR	32.04
436	Mar del Plata	ARG	35.33	485	Puducherry	IND	32.03
437	Guiyang	CHN	35.15	486	Dar es-Salaam	TZA	31.99
438	Pretoria	ZAF	35.10	487	Honghuagang	CHN	31.94

488	Vung Tau	VNM	31.81	537	Davao City	PHL	29.55
489	Hanoi	VNM	31.79	538	Sousse	TUN	29.50
490	Astrakhan	RUS	31.77	539	Tiexi	CHN	29.43
491	Posadas	ARG	31.67	540	Ciudad de Guatemala	GTM	29.41
492	Jinan	CHN	31.62	541	Changchun	CHN	29.25
493	Zhangzhou	CHN	31.60	542	Yaroslavl	RUS	29.22
494	Alajuela	CRI	31.53	543	Campinas	BRA	29.14
495	Tyumen	RUS	31.51	544	Lipetsk	RUS	29.10
496	Quito	ECU	31.50	545	Dongguan	CHN	29.08
497	Leshan	CHN	31.44	546	São José do Rio Preto	BRA	29.06
498	Tegucigalpa	HND	31.40	547	Novosibirsk	RUS	28.99
499	Fortaleza	BRA	31.38	548	Wuhu	CHN	28.99
500	Zhuzhou	CHN	31.37	549	Taiyuan	CHN	28.91
501	Baotou	CHN	31.35	550	Medellín	COL	28.87
502	Jiujiang	CHN	31.24	551	Campo Grande	BRA	28.84
503	Phuket	THA	31.24	552	Kaliningrad	RUS	28.81
504	Xishan	CHN	31.21	553	Eskisehir	TUR	28.80
505	Kaifeng	CHN	31.21	554	Cuiabá	BRA	28.67
506	Cusco	PER	31.11	555	Taguatinga	BRA	28.64
507	Nantong	CHN	31.06	556	La Plata	ARG	28.53
508	Almaty	KAZ	31.06	557	Panlong	CHN	28.52
509	Hefei	CHN	31.05	558	Mianyang	CHN	28.50
510	Ciudad Juárez	MEX	31.02	559	Irapuato	MEX	28.39
511	Magnitogorsk	RUS	30.98	560	Asunción	PRY	28.33
512	Cagayan de Oro	PHL	30.95	561	Toluca	MEX	28.13
513	Dujiangyan	CHN	30.90	562	Foshan	CHN	28.12
514	Chennai	IND	30.84	563	Skopje	MKD	28.11
515	San Juan	ARG	30.80	564	Bacolod	PHL	28.11
516	Rosario	ARG	30.80	565	Khabarovsk	RUS	28.10
517	Serrekunda	GMB	30.72	566	Goiânia	BRA	27.97
518	Joinville	BRA	30.71	567	Izmir	TUR	27.90
519	Manila	PHL	30.70	568	Belém	BRA	27.85
520	Tirana	ALB	30.67	569	Curitiba	BRA	27.84
521	Aguascalientes	MEX	30.62	570	Bukit Mertajam	MYS	27.82
522	Jodhpur	IND	30.62	571	San Luis Potosí	MEX	27.74
523	Kota Bharu	MYS	30.58	572	Chandigarh	IND	27.69
524	Beihai	CHN	30.55	573	Canoas	BRA	27.66
525	Penza	RUS	30.46	574	Niterói	BRA	27.66
526	Jinhua	CHN	30.28	575	San Salvador	SLV	27.65
527	Guadalajara	MEX	30.21	576	Medan	IDN	27.61
528	Córdoba	ARG	30.20	577	Zibo	CHN	27.60
529	Ribeirão Preto	BRA	30.12	578	Caxias do Sul	BRA	27.51
530	Vila Velha	BRA	30.06	579	Lomé	TGO	27.41
531	Tunis	TUN	29.98	580	Taishan District	CHN	27.39
532	Huadu	CHN	29.90	581	Vinnytsia	UKR	27.36
533	Cuenca	ECU	29.76	582	Yangon	MMR	27.26
534	Hyderabad	IND	29.63	583	Pereira	COL	27.23
535	Zigong	CHN	29.56	584	Belo Horizonte	BRA	27.23
536	Hohhot	CHN	29.56	585	Kajang	MYS	27.17

586	Mysuru	IND	27.16
587	João Pessoa	BRA	27.14
588	Yanji	CHN	27.06
589	Brest	BLR	27.03
590	Vientiane	LAO	26.92
591	Qingyuan	CHN	26.88
592	Villahermosa	MEX	26.81
593	Cheonan-si	KOR	26.78
594	Bursa	TUR	26.77
595	Bishkek	KGZ	26.77
596	Mahilyow	BLR	26.75
597	Zhengzhou	CHN	26.71
598	Santiago de Cali	COL	26.71
599	Belgorod	RUS	26.69
600	Bogotá	COL	26.67
601	Mangaluru	IND	26.67
602	Chiclayo	PER	26.65
603	Recife	BRA	26.62
604	Zhangjiakou	CHN	26.53
605	León	MEX	26.51
606	Pikine	SEN	26.47
607	Teresina	BRA	26.45
608	Porto Alegre	BRA	26.40
609	Can Tho	VNM	26.39
610	Pachuca	MEX	26.34
611	Baoding	CHN	26.21
612	Torreón	MEX	26.19
613	Ulaanbaatar	MNG	26.10
614	Jaboatão dos Guara.	BRA	26.00
615	Cebu City	PHL	25.98
616	Xinxiang	CHN	25.98
617	Bhubaneshwar	IND	25.97
618	Kharkiv	UKR	25.95
619	Irkutsk	RUS	25.94
620	Jundiaí	BRA	25.90
621	Quezon City	PHL	25.90
622	Iloilo City	PHL	25.80
623	Shantou	CHN	25.77
624	Yogyakarta	IDN	25.69
625	Tomsk	RUS	25.63
626	Durango	MEX	25.48
627	Uberlândia	BRA	25.38
628	Voronezh	RUS	25.29
629	Dandong	CHN	25.29
630	Hengyang	CHN	25.26
631	Manaus	BRA	25.22
632	Ibagué	COL	25.21
633	Surakarta	IDN	25.17
634	Beibei	CHN	25.17
635	Krasnoyarsk	RUS	25.05
636	Hermosillo	MEX	25.03
637	Indore	IND	25.02
638	Vitsebsk	BLR	25.00
639	Villavicencio	COL	24.93
640	Naberezhnye Chelny	RUS	24.88
641	Veracruz	MEX	24.80
642	Natal	BRA	24.65
643	Juiz de Fora	BRA	24.64
644	Kuching	MYS	24.57
645	Coatzacoalcos	MEX	24.55
646	Londrina	BRA	24.53
647	Saltillo	MEX	24.40
648	Da Nang	VNM	24.35
649	Yinchuan	CHN	24.33
650	Saratov	RUS	24.31
651	Culiacán	MEX	24.23
652	Seremban	MYS	23.99
653	Smolensk	RUS	23.76
654	Tucumán	ARG	23.65
655	São Luís	BRA	23.60
656	Kathmandu	NPL	23.45
657	Jingzhou	CHN	23.42
658	Cheboksary	RUS	23.32
659	Hengshui	CHN	23.27
660	Aracaju	BRA	23.24
661	Yongchuan	CHN	23.08
662	Arequipa	PER	23.05
663	Changping	CHN	22.85
664	Ufa	RUS	22.83
665	Langfang	CHN	22.81
666	Chelyabinsk	RUS	22.59
667	Klang	MYS	22.54
668	Hue	VNM	22.49
669	Ivanovo	RUS	22.40
670	São José dos Campos	BRA	22.38
671	Hunnan	CHN	22.31
672	Quilpué	CHL	22.25
673	Homyel	BLR	22.18
674	Maringá	BRA	22.05
675	Dehradun	IND	21.99
676	Novo Hamburgo	BRA	21.98
677	Bauru	BRA	21.97
678	Arusha	TZA	21.80
679	Agra	IND	21.71
680	Jilin	CHN	21.63
681	Mykolaiv	UKR	21.27
682	Ulan-Ude	RUS	21.22
683	Jakarta	IDN	21.19

684	Lima	PER	21.09
685	Ryazan	RUS	21.06
686	Ouagadougou	BFA	20.87
687	Kemerovo	RUS	20.84
688	La Paz	BOL	20.81
689	Uberaba	BRA	20.77
690	Batam	IDN	20.75
691	Ankara	TUR	20.72
692	Santiago de Cuba	CUB	20.69
693	Tepic	MEX	20.41
694	Riyadh	SAU	20.40
695	Semarang	IDN	20.15
696	Neiva	COL	20.07
697	Campina Grande	BRA	20.03
698	Pasto	COL	20.02
699	Mbour	SEN	20.01
700	Nagpur	IND	19.79
701	Shunyi	CHN	19.68
702	Stavropol	RUS	19.65
703	Mataram	IDN	19.53
704	Angeles	PHL	19.33
705	Orenburg	RUS	19.21
706	Cochabamba	BOL	19.11
707	Nashik	IND	19.00
708	Zaporizhzhia	UKR	18.95
709	Xalapa	MEX	18.62
710	Kisumu	KEN	18.37
711	Sorocaba	BRA	18.11
712	Tuxtla Gutiérrez	MEX	17.94
713	Manizales	COL	17.73
714	Guwahati	IND	17.58
715	Bucaramanga	COL	17.52
716	Barnaul	RUS	17.42
717	Luxor	EGY	16.79
718	Vadodara	IND	15.96
719	Lianyungang	CHN	15.72
720	Surabaya	IDN	14.53
721	Piracicaba	BRA	14.45
722	Trujillo	PER	13.99
723	Lucknow	IND	13.31
724	Villa Nueva	GTM	13.27
725	Bryansk	RUS	13.11
726	San Lorenzo	PRY	12.95
727	Managua	NIC	12.78
728	Bandung	IDN	10.95
729	Kumasi	GHA	10.79
730	Jeddah	SAU	10.72
731	Montería	COL	10.44
732	Srinagar	IND	10.28
733	Caracas	VEN	9.82
734	Mandalay	MMR	5.54

Table A2: Correlations, price of a representative apartment at the city center

	A	B	C	D	E	F	G	H	I	J	K	L	M
A	1.00												
B	0.96	1.00											
C	0.99	0.96	1.00										
D	0.99	0.96	0.99	1.00									
E	1.00	0.96	0.99	0.99	1.00								
F	0.97	0.94	0.97	0.96	0.97	1.00							
G	1.00	0.96	0.99	0.99	0.99	0.96	1.00						
H	1.00	0.96	0.99	0.99	1.00	0.97	1.00	1.00					
J	0.98	0.95	0.98	0.97	0.98	0.96	0.98	0.98	1.00				
K	0.95	0.93	0.95	0.94	0.95	0.93	0.95	0.95	0.97	1.00			
L	0.98	0.96	0.98	0.99	0.98	0.96	0.98	0.98	0.99	0.96	1.00		
M	1.00	0.96	0.99	0.99	0.99	0.97	0.99	1.00	0.97	0.95	0.98	1.00	
N	0.92	0.90	0.92	0.92	0.91	0.92	0.91	0.92	0.95	0.94	0.95	0.91	1.00

Note: For this table, I recompute the ranking of cities by their estimated nightly short-term rental rate of a representative apartment at the city center (Table A1). The table reports correlation coefficients of the estimated USD prices among all rankings. Specification A refers to the baseline version as shown in Table A1. B is based only on entire apartments, excluding properties that are shared. C controls for ratings. D controls for all possible amenities. E does not control for proximity to the shore of an ocean or big lake. F includes all properties that have been rented at least once. G includes all properties that have been rented or available at least 125 of 365 days. H uses prices not windsorized to the 0.01 and 0.99 percentiles by country. Specification I uses the centers from OSM for all cities. J uses the centers from Google Maps for all cities. K uses ln(distance) instead of ln(distance + 1) to compute the distance gradients. L divides the properties in each city in halfs, according to their distance to the city center, and then demeans by city halfs. Finally, M uses average retal prices instead of rental prices at the city center by not including distance gradients.

Table A3: Rank correlations, price of a representative apartment at the city center

	A	B	C	D	E	F	G	H	I	J	K	L	M
A	1.00												
B	0.94	1.00											
C	0.98	0.94	1.00										
D	0.99	0.94	0.97	1.00									
E	1.00	0.94	0.98	0.98	1.00								
F	0.93	0.89	0.93	0.92	0.93	1.00							
G	0.99	0.93	0.97	0.98	0.99	0.92	1.00						
H	1.00	0.94	0.98	0.99	1.00	0.93	0.99	1.00					
J	0.96	0.91	0.95	0.95	0.96	0.93	0.95	0.96	1.00				
K	0.92	0.89	0.92	0.91	0.92	0.90	0.91	0.92	0.95	1.00			
L	0.96	0.93	0.95	0.98	0.96	0.92	0.96	0.96	0.97	0.94	1.00		
M	0.99	0.93	0.98	0.98	0.99	0.93	0.99	0.99	0.95	0.92	0.95	1.00	
N	0.89	0.88	0.89	0.90	0.88	0.89	0.88	0.89	0.94	0.92	0.95	0.88	1.00

Note: For this table, I recompute the ranking of cities by their estimated nightly short-term rental rate of a representative apartment at the city center (Table A1). The table reports correlation coefficients of the estimated rank positions among all rankings. Specification A refers to the baseline version as shown in Table A1. B is based only on entire apartments, excluding properties that are shared. C controls for ratings. D controls for all possible amenities. E does not control for proximity to the shore of an ocean or big lake. F includes all properties that have been rented at least once. G includes all properties that have been rented or available at least 125 of 365 days. H uses prices not windsorized to the 0.01 and 0.99 percentiles by country. Specification I uses the centers from OSM for all cities. J uses the centers from Google Maps for all cities. K uses ln(distance) instead of ln(distance + 1) to compute the distance gradients. L divides the properties in each city in halfs, according to their distance to the city center, and then demeans by city halfs. Finally, M uses average retal prices instead of rental prices at the city center by not including distance gradients.

Figure A3: Comparison

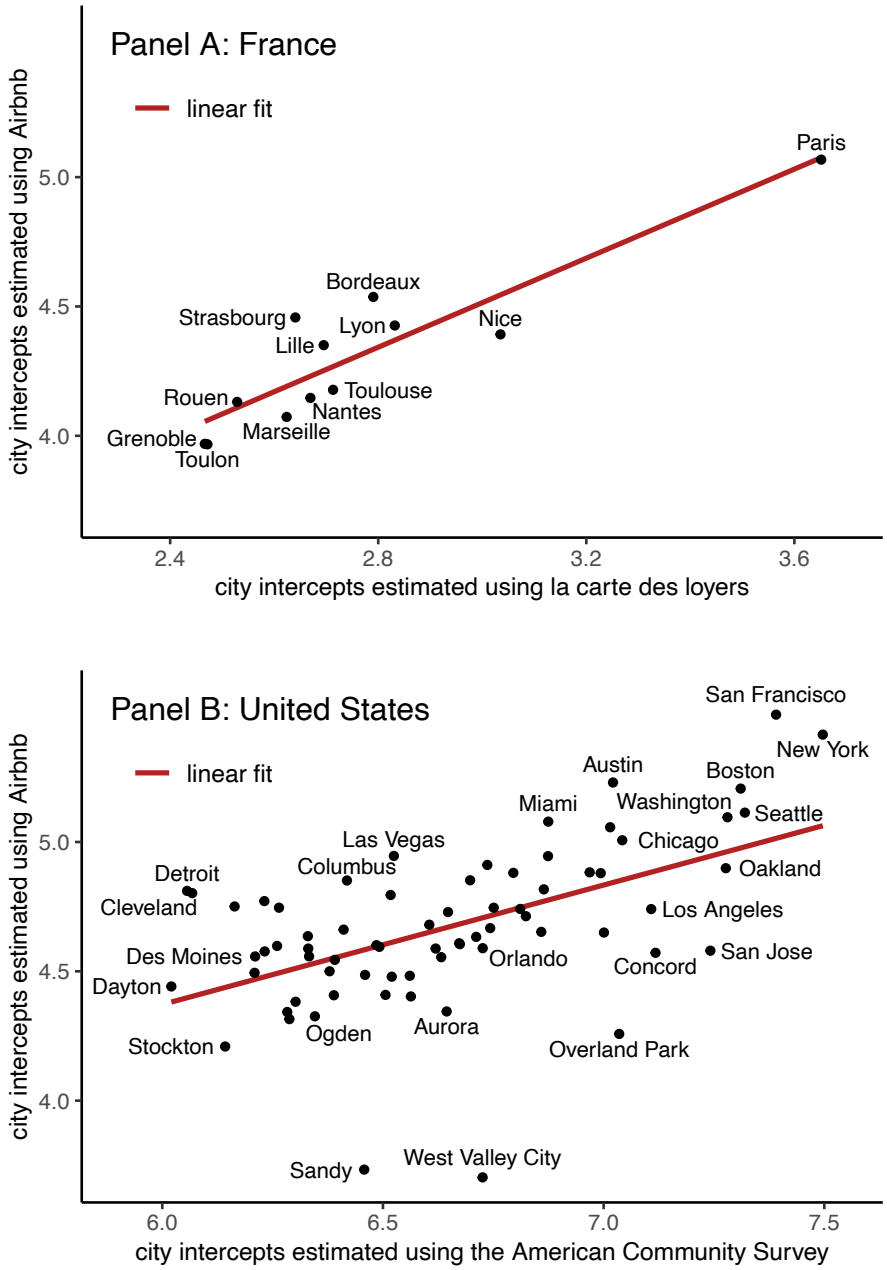

Note: This figure shows correlations between city intercepts estimated from short-term and long-term rental data. Panel A shows the analysis for France, using commune-level rents from *la carte des loyers*, who estimate them as outputs of hedonic regressions. I match 451 communes (and arrondissements) to my 12 (functional) French cities. I then regress prices on city distance gradients and city intercepts. The x-axis of Panel A shows these estimated intercepts. Panel B shows the analysis for the United States, using block group level rents from the 2015-2019 American community survey, which are averages of survey responses. I match 43,636 block groups to my (functional) US cities. I then regress prices on city distance gradients and city intercepts. In this case, I control for the block groups' fractions of several building-related characteristics. The x-axis of Panel B shows these estimates. In both cases, the y-axis shows the city intercepts that are the outcome of my first-stage regression using properties from Airbnb that lead to the ranking in Table A1.

Figure A4: Compactness

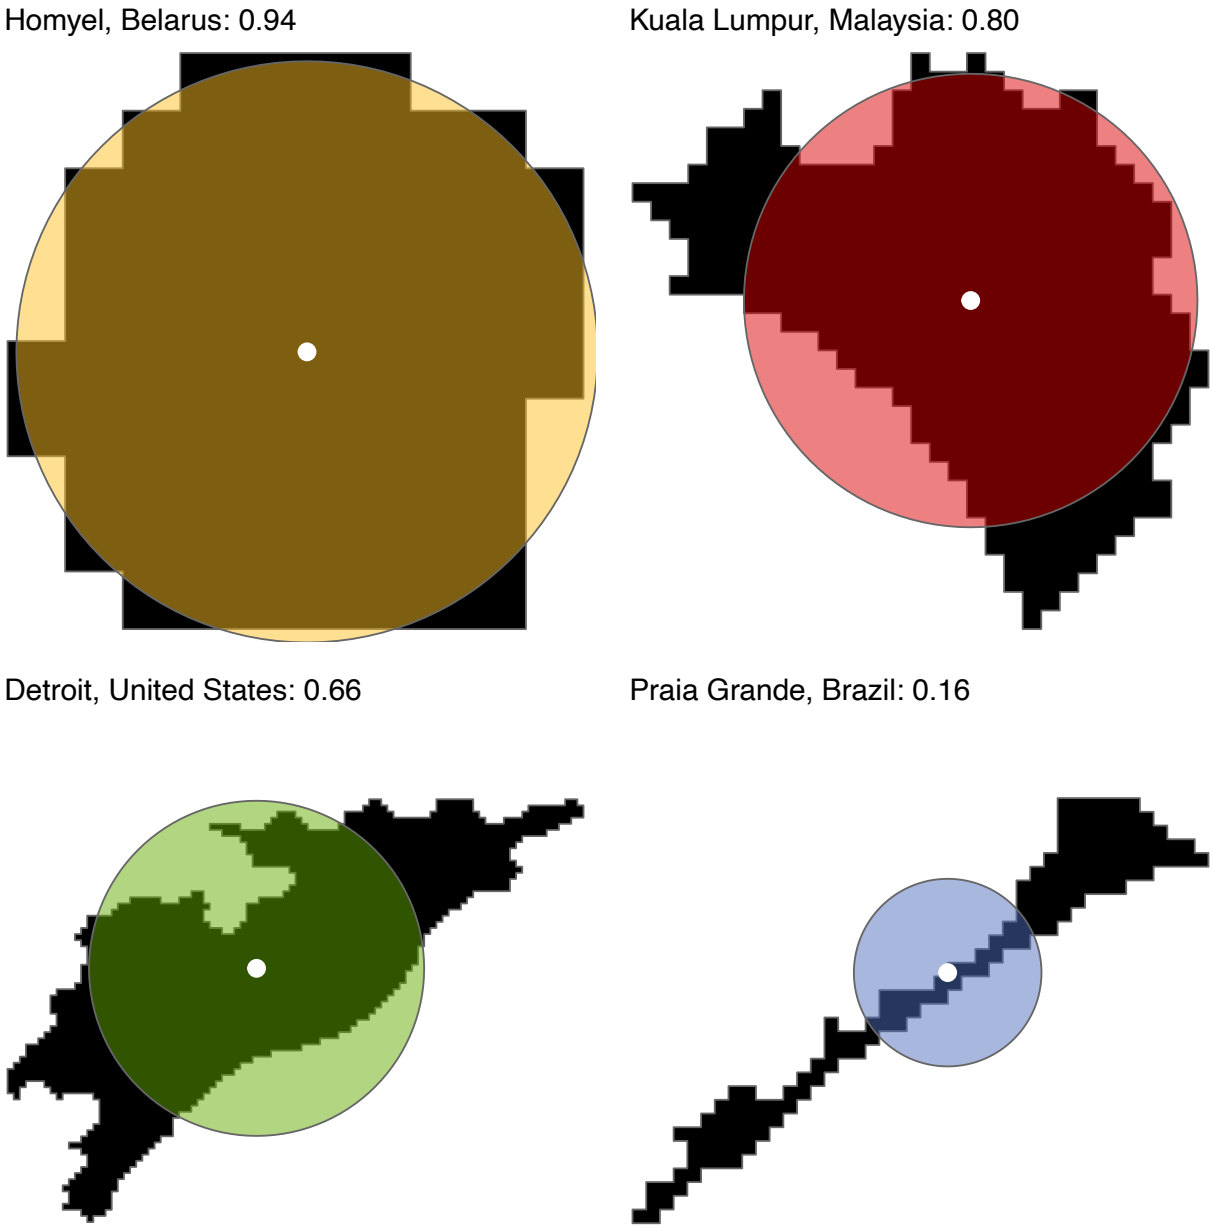

Note: The figure shows the compactness measure that I use as a second-stage control variable for four exemplary cities. The chosen cities have compactness measures that correspond to the maximal value, 75% quantile, 25% quantile, and minimal value of the distribution. The measure is taken from Angel et al. (2020) where it is called "exchange". To create it, I compute a circle with the same area as the city itself around each city's centroid and then measure the proportion of the circle that intersects with the shape of the city.

Table A4: Additional specifications

Dependent Variable:	log(Price of representative apartment at city center)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
log(Population)	-0.435***		-0.330***		0.164***	0.161***
	(0.062)		(0.083)		(0.057)	(0.050)
log(Area)	0.639***		0.516***		-0.030	-0.050
	(0.068)		(0.104)		(0.063)	(0.053)
Compactness		-0.198	0.078			-0.054
		(0.313)	(0.195)			(0.112)
Elevation (100m)		-0.014***	-0.004			-0.007*
		(0.003)	(0.004)			(0.004)
Difference to 21.11°C		0.010	0.004			0.002
		(0.014)	(0.012)			(0.007)
By ocean / big lake		0.111**	0.100***			0.054
		(0.051)	(0.035)			(0.045)
Capital		0.123	0.114**			0.105*
		(0.100)	(0.055)			(0.057)
Airbnbs per 1,000		0.058***	0.048***			0.029***
		(0.010)	(0.009)			(0.006)
Country fixed effects	-	-	-	Yes	Yes	Yes
R^2	0.325	0.281	0.475	0.736	0.776	0.813

Note: The table shows regressions of the estimated price of a representative short-term rental property at the city center on city size and control variables. The units of observation are 733 cities. Column 4 includes only country fixed effects. The parentheses show standard errors, which are clustered by country. The levels of significance are * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Heterogeneity by country without the countries' two largest cities

Dependent Variable:	log(Price of representative apartment at city center)						
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	USA	Eurozone	Russia	China	India	Brazil	Mexico
log(Population)	0.23**	0.30***	0.02	0.10	0.42*	0.04	-0.45**
	(0.11)	(0.09)	(0.26)	(0.12)	(0.22)	(0.09)	(0.19)
log(Area)	-0.07	-0.10	0.14	0.01	-0.23	0.04	0.59***
	(0.12)	(0.09)	(0.26)	(0.13)	(0.23)	(0.11)	(0.19)
Compactness	-0.36**	0.35	-0.34	-0.34	-0.17	0.34	-0.65
	(0.18)	(0.36)	(0.43)	(0.32)	(1.30)	(0.24)	(0.54)
Elevation (100m)	-0.02***	-0.01	-0.04*	0.00	0.00	0.00	0.00
	(0.01)	(0.03)	(0.02)	(0.01)	(0.03)	(0.01)	(0.01)
Difference to 21.11°C	0.02**	0.08***	0.00	-0.01**	-0.05	0.00	-0.03
	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)	(0.01)	(0.02)
By ocean / big lake	-0.11**	-0.03	0.04	0.22***	0.15	0.01	0.07
	(0.05)	(0.07)	(0.07)	(0.08)	(0.29)	(0.07)	(0.14)
Airbnbs per 1,000	0.04***	0.02***	0.06**	0.03**	0.78**	0.10***	0.10***
	(0.01)	(0.00)	(0.03)	(0.01)	(0.36)	(0.03)	(0.02)
Country fixed effects	-	Yes	-	-	-	-	-
Observations	68	74	42	110	29	42	36

Note: The table shows regressions of the estimated price of a representative short-term rental property at the city center on city size and control variables. The units of observation are cities. The two cities with the largest population in each entity are excluded. The parentheses show standard errors clustered by country for specification 2 (eurozone) and heteroscedasticiy robust standard errors for all other specifications. The levels of significance are * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Comparison with Chauvin et al. (2017)

		USA		Brazil	China		India
		ln(rent)	ln(price)	ln(rent)	ln(rent)	ln(price)	ln(rent)
Chauvin et al. (2017)	OLS	0.15***	0.20***	0.13***	0.23***	0.10	0.003
		(0.01)	(0.04)	(0.02)	(0.08)	(0.12)	(0.005)
	IV	0.15***	0.20***	0.13***	0.37***	0.06	−0.004
		(0.01)	(0.04)	(0.02)	(0.13)	(0.13)	(0.009)
		log(nightly rate of representative apartment at city center)					
Preferred	OLS	0.17***		0.09***	0.13***		0.23***
		(0.03)		(0.02)	(0.03)		(0.08)
	IV	0.18***		0.08***	0.13***		0.23**
		(0.03)		(0.02)	(0.03)		(0.09)
No controls	OLS	0.19***		0.11***	0.14***		0.16*
		(0.04)		(0.03)	(0.03)		(0.08)
	IV	0.20***		0.10***	0.11***		0.16*
		(0.04)		(0.03)	(0.04)		(0.08)
		log(nightly rate of representative apartment anywhere in the city)					
No controls, means	OLS	0.12***		0.09***	0.06**		0.03
		(0.03)		(0.03)	(0.02)		(0.04)
	IV	0.12***		0.10***	0.02		0.02
		(0.03)		(0.03)	(0.03)		(0.04)

Note: The first strip shows the estimates Chauvin et al. (2017) obtain when regressing either log rents or log house prices on log population. They measure population in 2010 and use population in 1980 as an instrument. The other three stripes are based on my own estimates. The dependent variable of the second and the third strip is the estimated price of a representative short-term rental property at the city center. The dependent variable of the fourth strip is the estimated price of a representative short-term rental property anywhere in the city. The difference between the two is whether the first-stage hedonic regression does (strip 2 and 3) or does not (strip 4) include distance gradients. The second strip is estimated using all control variables that are included in Table 3, but does not control for area. The third and the fourth strip are estimated using neither controls nor controlling for area. Population in 2015 is instrumented by population in 1975 for the IV specifications of stripes 2, 3, and 4. The parentheses show heteroscedasticity robust standard errors. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Chapter 3

Trade and conflict in Myanmar: A reverse China shock

Trade and conflict in Myanmar: A reverse China shock[1]

Bernhard Nöbauer[2]
University of Lausanne

September 2023

Abstract

I study the effect of trade in mining goods on conflict events in the border region of Myanmar. Using a shift-share measure, I disaggregate national exports to the township level. Imports from other low and middle-income countries are used to construct an instrumental variable to rule out reverse causality. I use a two-way fixed effects model for estimation. Export exposure to mining goods is associated with an increase in violent conflict. This increase predominantly affects townships inhabited by ethnic minorities. A placebo test confirms that the production of mining goods drives the effect. Night lights close to the mines are brighter in years with high export exposure, but surrounding areas do not seem to benefit.

[1] I am very thankful to Marius Brülhart for his invaluable guidance. I thank Gilles Duranton, Mariaflavia Harari, Jeremias Kläui, Dzhamilya Nigmatulina, Dominic Rohner, Mathias Thoenig, and the participants of the Gerzensee Alumni Conference 2022 for helpful comments. Moreover, I thank Mélise Jaud, Henri Rueff, and Kevin Woods for their insights on the situation in Myanmar.

[2] Departement of Economics, Faculty of Business and Economics (HEC Lausanne), University Of Lausanne, 1015 Lausanne, Switzerland; `bernhard.noebauer@bluewin.ch`.

1 Introduction

Do increased exports of natural resources fuel conflict? This paper looks at Myanmar's[1] exports of mining goods to China over the period 2012 to 2020. It asks whether they can help to explain the high level of conflict in Kachin and Shan, the two states of Myanmar that share a border with China. I use a Bartik (1991) style shift-share measure to quantify a township's export exposure to mining goods. First, I disaggregate nationwide export values in different metals to the townships hosting mines associated with these metals. Second, I construct an instrumental variable (IV) to deal with endogeneity concerns, using Chinese imports from other low and middle-income countries as an input. Effects are assessed using a two-way fixed effects model. Export exposure in mining goods increases the number of conflict events, with an elasticity of 0.46 for the OLS specification. This estimate is likely biased towards zero due to reverse causality, as higher levels of local conflict hamper exports. The elasticity increases to around 0.54 when using the IV approach. Ethnicity seems to matter, as most of the increase in conflict events is due to townships not inhabited by the nationwide ethnic majority. A placebo test further validates the results, replacing trade flows in mining goods with those in other goods of comparable importance for the export sector, typically produced elsewhere. As expected, the effect disappears when using this counterfactual measure of export exposure in placebo goods. Furthermore, I analyze the effect of export exposure on night lights as a proxy of economic development. In the immediate neighborhood of mines, pixels are brighter in years with higher export exposure. However, the effect does not spill over to nearby areas. Areas inhabited by ethnic minorities experience a lower increase in night lights when export exposure is higher, with the effect disappearing even quicker with distance.

When it comes to the effect of trade on civil conflict, the evidence is quite limited. Martin et al. (2008) find that international trade decreases the risk of severe civil wars, as the potential halt of this trade constitutes an additional cost. However, it might make low-scale conflict more likely, as trade within countries is less crucial if international alternatives exist. Brülhart et al. (2019) look at cross-border trade and find that for Africa an increase in cross-border trade is associated with a reduction in conflict. Candau et al. (2022) find that trade decreases the probability of ethnic wars in Africa, although this effect is driven mainly by agricultural products, while mining goods show no significant effect. The research by Calì and Miaari (2015) is probably closest to my work. They use a shift-share set up to assess the effect of Palestinian exports (mainly to Israel) on violence during the Second Palestinian Intifada. They argue that the Palestinian economy does not feature any extractive industries and that the state capacity channel is negligible, leading them to estimate an isolated opportunity cost channel (p.4). Consequently, they find a pacifying effect of export exposure. In contrast, in the case of Myanmar, the opportunity cost channel (better outside options in the formal sector), the rapacity channel (bigger prize that can be obtained by fighting), and the state capacity channel (higher income allows more funding of government forces) all play a role, which implies that the direction of the result is a priori not evident. Moreover, Calì and Miaari (2015) only consider a cross-section, while my data allows me to exploit a panel structure, controlling for constant factors across townships or

[1] No deliberate stance is implied by the usage of the word Myanmar instead of Burma in this article.

years. There are also studies considering the effect of war on trade. Glick and Taylor (2010) quantify a large negative impact of war on trade, while Qureshi (2013) shows that war has spillover effects, decreasing trade for neighboring countries, even if they do not participate in the conflict. The fact that there are papers highlighting both directions of causality calls for a thorough identification strategy.

While the evidence regarding the effect of trade on conflict is not unambiguous, it generally points towards a pacifying effect of trade. At the same time, there is a substantial literature relating world market prices of mining goods to conflict. This literature typically reports higher levels of conflict when the value of minerals rises (see, for example, Berman et al., 2017). As argued by Dube and Vargas (2013) and Dal Bó and Dal Bó (2011), a plausible explanation for this finding is that mining goods are less labor intensive than other goods, dampening the effect of the opportunity cost channel that causes workers to participate in the formal sector when prices are high. At the same time, mining goods are often spatially concentrated and more easily lootable than other goods (Ferraz et al., 2021).

There are good reasons to assume that the effect of trade goes beyond a mere price effect. The actual occurrence of trade implies that trade links have been established, contracts have been signed, workers have been employed, and the metals have been extracted and transported to the border. All of these steps require the cooperation of many different people from potentially different ethnic backgrounds, but they might also give rise to grievances. Candau et al. (2022, p.532) argue that "[i]n analysing global exogenous price shocks, many papers focus on the short-run effects of international trade." At the same time, focusing on actual trade faces the problem of reverse causality. In this regard, Myanmar provides an ideal laboratory, as most trade occurs with China, a trading partner which is far bigger and whose demand for natural resources is arguably independent of events within Myanmar. Despite the focus on trade, I confirm the positive effect of mining on conflict often reported in the literature.

My work contributes to the class of within-country studies on natural resources and conflict that is still limited but growing in recent years (see, for example, Aragón and Rud, 2016; Crost and Felter, 2020). These studies have the advantage that the institutional environment is relatively similar across the study area. Furthermore, the panel structure enables me to control for factors that are specific to a given township or year but invariant over the other dimension. Novel data allow me to work on the township level, the lowest level covered by the conflict dataset. Myanmar is a multi-ethnicity state, with the ethnic majority being clustered predominantly in the center of the country and controlling the government. This setting is not unusual in the literature, and ethnic composition has been used to explain the differential effects of natural resource wealth on conflicts (Janus and Riera-Crichton 2015, Morelli and Rohner 2015, Giménez-Gómez and Zergawu 2018). However, Myanmar differs from other settings because of the strength of the military and the level of repression it can exert. Arezki and Brueckner (2021) find that countries with high military expenditures do not exhibit a positive relation between rents from natural resources and civil conflict. In this context, the state capacity channel could explain why townships populated by, among others, the ethnic majority do not experience an export exposure effect on conflict.

Moreover, the paper sheds light on an effect of China's unprecedented rise over the last decades. By the sheer size of the development, it significantly affects many other parts of the global economy. Some of these consequences are well researched. Autor et al. (2013) find that local labor markets in the United States that experienced higher import exposure to Chinese manufacturing goods faced increased unemployment and reduced wages. Their seminal work inspired many other papers examining various plausible consequences, such as on worker health (McManus and Schaur, 2016) and voting behavior (Colantone and Stanig 2018, Autor et al. 2020). Dauth et al. (2017) also look at export exposure to China rather than just import exposure from China, finding that export exposure to China has slowed the decline of the German manufacturing sector. However, to my knowledge, most of the literature using this identification strategy to document the consequences of the rapid trade growth with China is concerned with effects on high-income countries. This focus is surprising, as trade with low and middle-income countries has also proliferated. Standard trade theory would suggest that the implications of this increase will be different for low and middle-income countries. Rather than looking at imports of manufactured goods from China, this paper considers exports of primary sector mining goods to China, driven by China's increasing demand for natural resources. In doing so, it mirrors the estimation strategy of Autor et al. (2013). While they use imports of other high-income countries from China to construct their instrument, this work considers Chinese imports from other low and middle-income countries.

This paper proceeds as follows: Section 2 explains the context in Myanmar and presents the data. Section 3 outlines the estimation method. Section 4 describes the results and explores their robustness. It also includes an analysis of local spillover effects using night lights. Section 5 provides a discussion and concludes.

2 Empirical Setting

This paper combines data from multiple sources, all of which are publicly available. The study period is defined by the availability of these data. The trade data used for this paper are currently available until 2020, while the night light data start in 2012.[2] This nine-year period was an eventful one in the history of Myanmar. It contains a slow process of democratic opening, with the National League for Democracy winning large majorities in general elections in 2015 and 2020. At the same time, the military still held considerable power, both in the democratic process and in the economy. The study period ends shortly before the coup d'état in February 2021.

The geographic focus of this paper is on Kachin and Shan, the two states within Myanmar that share a border with China (see Figure 2a). Nationwide, the largest ethnic group are the Bamar, who constitute about two-thirds of the population of Myanmar. They speak Burmese as their mother tongue, are predominantly Theravada Buddhist, and hold most of the positions

[2] Night light data also exist for the years before 2012. However, there was a discontinuity due to a change of satellites in 2012. There is no agreed on way to reconcile the data from the two different satellite systems (Chen and Nordhaus, 2015). Elvidge et al. (2013) provide a discussion on the many advantages the VIIRS system offers relative to the old DMSP-OLS system.

in the central government. Their settlement area also covers parts of Kachin and Shan (see Figure A2). However, these states are populated to a large extent by other ethnic groups like the Kachin, Shan, Wa, or Kayin. Both states host rebel groups. These groups often represent an ethnic minority and are referred to as ethnic armed organizations in these cases, a prominent example being the Kachin Independent Army. The military of Myanmar, the Tatmadaw, only partially asserts control over Kachin and Shan. In some places they do not fully control, they collaborate with paramilitary organizations, some of which are former rebel groups, that involve rent-sharing of the region's many natural resources. This situation regularly results in violent conflict between the different armed groups. Moreover, there are frequent cases involving violence against civilians, predominantly but not exclusively, committed by government forces (source: ACLED). For a detailed account of the situation and the different parties, see, e.g., Woods (2018).

2.1 Trade

The *Centre d'Études Prospectives et d'Informations Internationales* (CEPII) provides yearly trade data in the form of the *Base pour l'Analyse du Commerce International* (BACI).[3] Over the study period, trade between Myanmar and China increased sharply. Trade volume from and to China more than doubled from 7.5 billion USD in 2012 to almost 17 billion USD in 2020. As of 2020, China is by far the biggest trading partner for Myanmar, ahead of Thailand (6.6 billion USD), Singapore (2.8), and Japan (2.1) (BACI).

Figure 1: Development of trade in primary sector metals between 2012 and 2020.

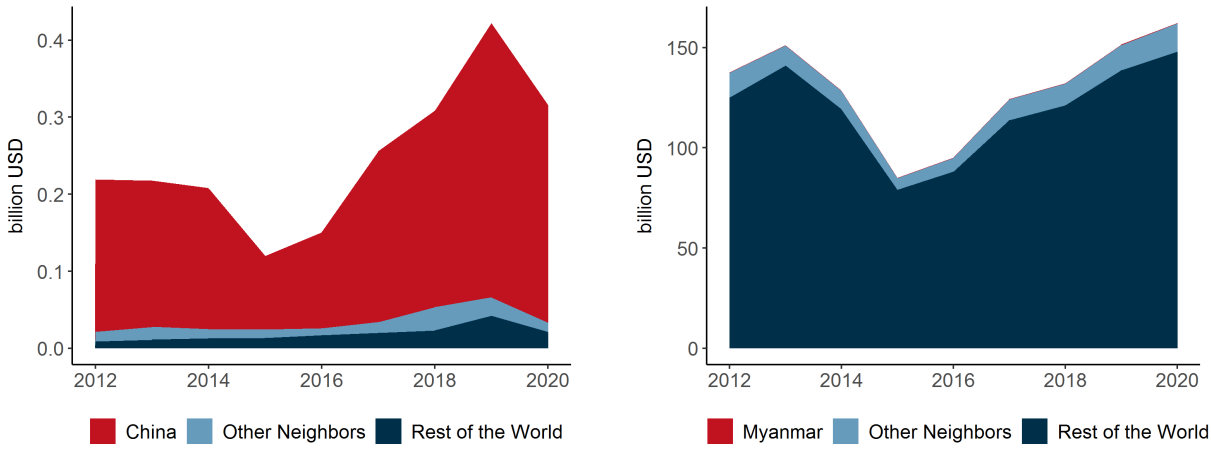

(a) Value of metals exported from Myanmar by destination.

(b) Value of metals imported by China by origin.

This dominance is even more striking considering the trade of primary sector mining goods, which is the focus of this paper.[4] Figure 1a shows the evolution of exports from Myanmar for this set of goods. From 2012 to 2020, almost 86% of the value obtained by exporting mining

[3] Downloaded on 01.12.2022 from `http://www.cepii.fr/CEPII/en/bdd_modele/presentation.asp?id=37`.

[4] Table A1 in the Appendix shows how different goods are mapped to the metals of interest for this analysis. Mining goods are included if i) they are part of the primary sector as defined by the Broad Economic Categories Rev 4., ii) they can be linked to a mine in Kachin or Shan state, and iii) BACI reports exports from Myanmar to China for at least one year in the study period. Figures 1a and 1b refer to this subset of mining goods.

goods was due to trade with China. During this period, there is an unmistakable resemblance between exports from Myanmar and the development of Chinese imports, which are depicted in Figure 1b. However, Myanmar only accounted for a tiny share of Chinese imports in mining goods, about 0.2%, from 2012 to 2020. Therefore, it seems safe to assume that variations in Chinese demand have a substantial impact on the mining sector of Myanmar while Chinese demand for mining goods is at most marginally affected by supply shocks within Myanmar.

2.2 Mining

The mining data used here are based on the 6th Extractive Industries Transparency Initiative (EITI) report, covering the fiscal year 2018 (MEITI, 2020).[5] The data have been processed and geocoded by the Myanmar Centre for Responsible Business (MCRB). Apart from mine coordinates, the dataset also includes information about mine size, type, and the material exploited (MCRB, 2022) To my knowledge, these data have not yet been used in an academic context.

There were 1,218 active mining licenses in the fiscal year 2018 in Myanmar, 399 (32.8%) of which were located in Kachin and Shan states. Figure 2c shows the spatial distribution of these mines. The most frequent mine types are lead & zinc (110 mines), followed by coal (62), antimony (41), and tin & tungsten (29).[6]

The EITI report also includes data about jade and gems. Jade is significant for the economy of Myanmar, accounting for anywhere between 2% (MEITI, 2015) and 48% (Global Witness, 2015) of GDP (Oak, 2018). The sector is notoriously intransparent as the range of these estimates demonstrates. Moreover, jade production is heavily clustered. Five townships account for all production in jade and gemstones, three of which are in Kachin and Shan.[7] Finally, Myanmar is the biggest jade producer globally. As a consequence, developments in Myanmar inevitably influence jade trade worldwide. There is no straightforward way to use variation that is exogenous to Myanmar to construct a credible instrumental variable. Export exposure in jade and gems will therefore only be included as a control variable.

2.3 Conflict

Data about conflict events come from the Armed Conflict Location & Event Data Project (ACLED, Raleigh et al. 2010).[8] Figure 2a depicts the number of conflict events during this period, aggregated by state. With the exemption of Rakhine state, home of the Rohingya,

[5] In 2014, Myanmar became a member of the EITI, a global standard through which countries commit to publicly disclosing information about the extraction of their natural resources. Its membership is currently suspended as a consequence of the military coup (`https://eiti.org/countries/myanmar`, last accessed: 20.09.2022).

[6] The complete list of mines in Kachin and Shan by the material (mainly) extracted there includes: lead & zinc (110 mines), coal (62), antimony (41), tin & tungsten (29), iron (17), manganese dioxide (14), gypsum (11), marble (11), limestone (7), copper (6), baryte (4), quartzite (4), chromium (2), and dolomite (1). Moreover there are mines extracting gold (79) and bauxite (1). As there are no reported exports to China for the latter two metals, they are excluded from the study and not depicted on the map.

[7] In the trade data, it is not possible to distinguish jade from other gemstones. Global Witness (2015) argues that almost all of the export value is due to jade (p.101 f.). In the rest of the paper, I will just refer to this category as jade, implicitly including other types of gemstones.

[8] There exists also the Uppsala Conflict Data Program Georeferenced Event Dataset, which serves a similar purpose. However, for the case of Myanmar, ACLED has more than 9 times as many entries over the study period. For that reason, I prefer to work with ACLED for this analysis.

Figure 2: Geography of Myanmar (Panel a), with a particular focus on Kachin and Shan states (Panels b, c, and d).

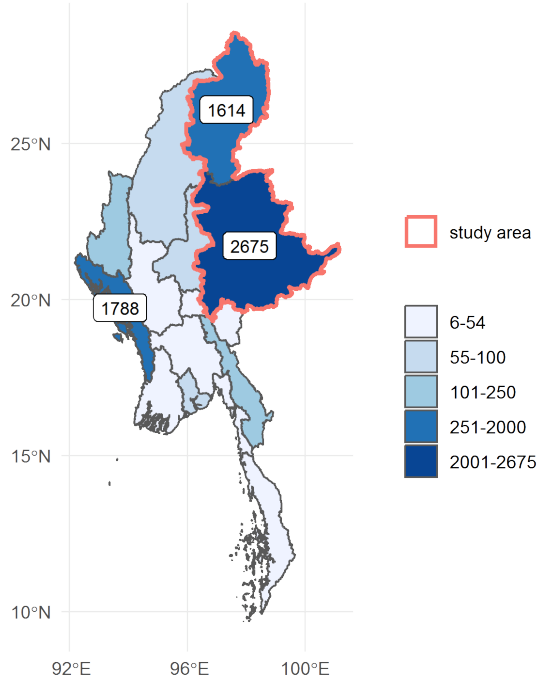

(a) Number of violent conflict events between 2012 and 2020 by state.

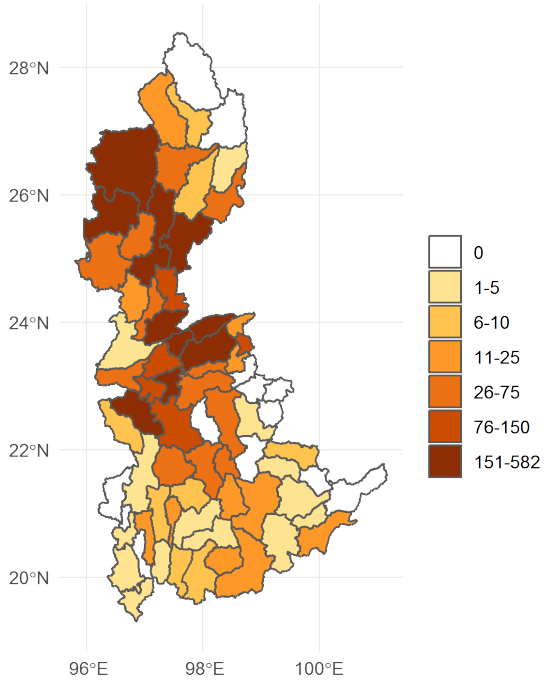

(b) Number of violent conflict events between 2012 and 2020 by township in Kachin and Shan.

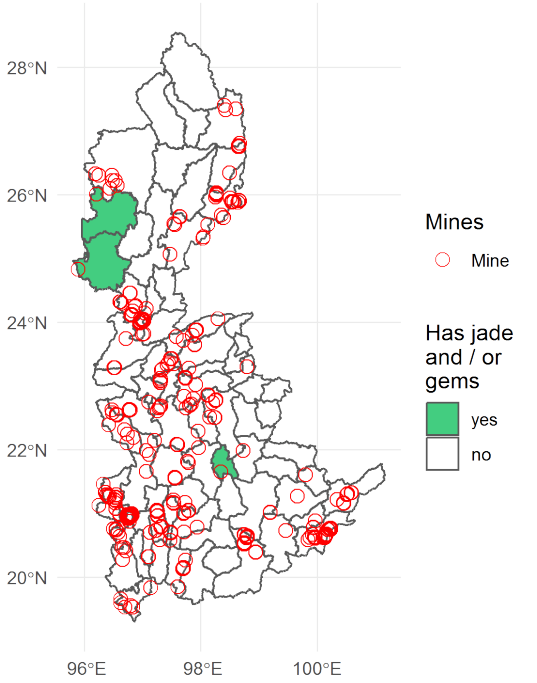

(c) Active mining licenses in the fiscal year 2017/18 in Kachin and Shan states.

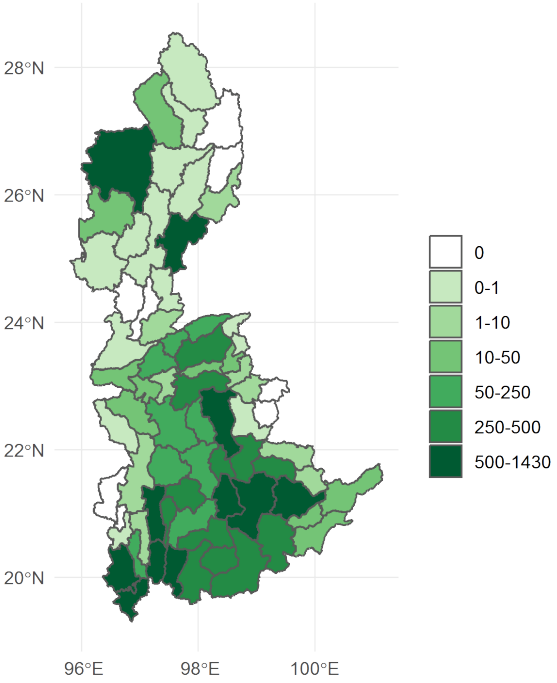

(d) Area of opium cultivation in hectares by township in Kachin and Shan states.

Kachin and Shan are by far the most conflict-ridden states in Myanmar. Figure 2b shows the spatial distribution of conflict events within Kachin and Shan by township. Figure A1 provides corresponding maps with conflict fatalities.

2.4 Opium

Myanmar is estimated to be the second largest opium producer in the world, behind Afghanistan (Kramer, 2017). Over 96% of the opium cultivated in Myanmar is grown in Kachin and Shan states (UNODC, 2021, p. 8). Moreover, the value of the produced opium has varied widely over the years. It seems, therefore, imperative to include a measure of the value of opium produced in a township in any given year. However, given the illegality of the sector, data are scarce and inevitably imprecise. This paper makes use of information provided by the United Nations Office on Drugs and Crime (UNODC), which produces yearly reports about opium cultivation in Myanmar. Appendix B describes the steps I take to convert this information into a variable measuring export exposure in opium by township and year.

2.5 Night lights

Night light data are a good proxy for economic activity and development levels, especially in countries where geographically disaggregated statistics about these factors are unreliable (Henderson et al., 2012). Data are collected by the "Visible and Infrared Imaging Suite (VIIRS) Day Night Band (DNB)" (Elvidge et al., 2021) and made available by the Earth Observation Group of the Payne Institute for Public Policy.[9]

2.6 Other data

Township and state boundaries are taken from the Myanmar Information Management Unit. Data on the presence of different ethnic groups come from the Ethnic Power Relations data sets (Vogt et al., 2015). Figure A2 depicts the settlement area of the ethnic majority (the Bamar) according to these data.[10] The correspondence table used to determine which goods belong to the primary sector is provided by the United Nations Statistics Division. Finally, the classification of low and middle-income countries is taken from the World Bank.

3 Estimation

To estimate the effect of trade exposure on conflict, I use a two-way fixed effects model, controlling for township and year fixed effects. To control for potential endogeneity, I employ a Bartik-style shift-share instrument.

I construct a measure for the exposure of a township to mining exports to China, which will be my variable of interest. For each metal m, the nationwide export value for this metal in year t is disaggregated to all townships i in which the metal is mined. However, mineral sources vary considerably in size. In my main specification, I do the disaggregation with regard to the

[9] More precisely, average values from the "Annual VNL V2" are used.
[10] Parts of this area overlap with the settlement areas of other ethnic groups.

share a township has in the nationwide mining area of the metal m. If, for instance, a township possesses 5% of the nationwide mining area of iron, I attribute 5% of the export value of iron for each year to this township. As a township can be endowed with several different metals, I then aggregate these exposures.

$$\text{Export Exposure}_{it} = \sum_{m} \text{Share}_{im} \times \text{Exports}_{mt}^{\text{MMR}\rightarrow\text{CHN}} \tag{1}$$

Regressing a measure of conflict intensity, for example the number of conflict events, on this measure of export exposure is the basis of what can be seen as the ordinary least squares (OLS) version of the analysis. Townships vary from one another in their proneness to conflict due to many different factors not related to trade. All regressions include township fixed effects λ_i to control for all such time-invariant factors. Moreover, they include year fixed effects μ_t to account for the fact that the overall situation in Kachin and Shan was more fragile in some years and more stable in others. I estimate the following OLS specification in which X_{it} consists of control variables for export exposure in jade and opium.

$$\text{Conflict Events}_{it} = \alpha + \beta \times \text{Export Exposure}_{it} + \gamma X_{it} + \lambda_i + \mu_t + \varepsilon_{it} \tag{2}$$

The two fixed effects attenuate the worry that the regression will be biased due to omitted variables. However, the analysis might also suffer from reverse causality as we might expect conflict to obstruct trade. As exports of a township are inferred by allocating a fraction of overall exports on a national scale even the OLS specification has a certain level of robustness against this problem. To go one step further, I use Chinese mining imports from other low and middle-income countries (LMIC) to construct a shift-share instrument of the following form.

$$\text{Export Exposure}_{it}^{\text{IV}} = \sum_{m} \text{Share}_{im} \times \text{Exports}_{mt}^{\text{LMIC}\rightarrow\text{CHN}} \tag{3}$$

Doing this yields a measure of counterfactual export exposure in mining goods for each township. It captures the part of the increase in trade driven by China's rising demand for natural resources. As Myanmar only accounts for a small share of overall Chinese imports, this measure should not be influenced by conflict within Myanmar. The IV version of the analysis can be expressed as

$$\text{Export Exposure}_{it} = a + b \times \text{Export Exposure}_{it}^{\text{IV}} + cX_{it} + l_i + m_t + e_{it} \tag{4}$$

$$\text{Conflict Events}_{it} = \alpha + \beta \times \widehat{\text{Export Exposure}}_{it} + \gamma X_{it} + \lambda_i + \mu_t + \varepsilon_{it} \tag{5}$$

4 Results

4.1 Main specification

Table 1 shows the regression results of the baseline specifications. All regressions are estimated with two-way fixed effects (73 townships and 9 years) and with two-way clustered standard errors (at the township and year level). I take logs off all variables so the coefficients can be interpreted as elasticities. There is a positive relationship between export exposure in metals and the number of conflict events. Columns (1) and (3), as well as (2) and (4), show very similar coefficients of interest. This implies that controlling for export exposure in jade and opium does not make a big difference. The OLS specifications show an elasticity of around 0.44. The elasticity increases to around 0.54 with the IV specification. This difference is consistent with (some degree of) reverse causality, causing us to underestimate the true effect when using simple OLS.

Table 1: Main specification

Dependent Variable:	ln(conflict events)					
Model:	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV
ln(export exposure metals)	0.42**	0.52*	0.46**	0.56*	0.72**	1.07***
	(0.17)	(0.25)	(0.17)	(0.26)	(0.22)	(0.30)
Bamar × ln(export exposure metals)					-0.75**	-1.55**
					(0.27)	(0.60)
ln(export exposure jade)			0.15	0.14	0.14	0.12
			(0.20)	(0.21)	(0.20)	(0.20)
ln(export exposure opium)			-0.43	-0.44	-0.31	-0.18
			(0.26)	(0.26)	(0.28)	(0.32)
Cragg-Donald Wald F statistic		206.74		203.16		98.49

Note: The table shows regressions of violent conflict events on export exposure in primary sector mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. I obtain export exposure by disaggregating nationwide exports to the townships where the goods are produced, using the townships' share of the national production area. The export exposure in jade and in opium is computed in the same way. Specifications five and six consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

In columns (5) and (6), export exposure in metals is interacted with an indicator variable for the presence of the Bamar ethnic majority. The positive effect of export exposure in mining goods on conflict seems to be driven entirely by townships that are completely populated by ethnic minorities. For townships in which the ethnic majority is present, the effect becomes very small in the OLS specification and negative in the IV specification.

I construct the export exposure in jade and opium in a similar way to the export exposure in metals. In all cases, exports on a national level are disaggregated to individual townships according to a township's share of national production.[11] Jade exposure also shows a consistently positive relation with conflict events, while export exposure in opium is negatively associated with conflict in most specifications. However, for both control variables, the coefficients are far from being statistically significant. The data on jade and opium seem to be particularly imprecise. For example, MEITI reports individual licenses for other metals, but only township aggregates for jade. Measurement error might therefore explain part of the insignificance when it comes to the control variables.

4.2 Robustness

Control variables

Table A2 shows that the specific set of control variables is not critical for the results. Even if the inclusion of jade exposure does not considerably alter the results, townships involved in the jade sector could greatly influence them. To attenuate this worry, Panel A excludes the three townships for which MEITI reports the presence of mines extracting jade. Panel B only controls for export exposure in jade, but not opium.

Allocating exports to townships

Table A3 explores the robustness to the allocation of nationwide exports to townships. In the main specification, for each material, exports are disaggregated by the share of a township's mining area relative to the whole country. There are a few licenses linked to exploration for which unrealistically large values are reported. Before the disaggregation is done, I set these values to the largest area of a mine that is classified as producing. Panel A of Table A3 shows the results without this adjustment.

Panel B proposes a disaggregation method that is more robust to outliers and erroneous information about the size of a mining area. MEITI categorizes mines into four categories: large-scale production, small-scale production, subsistence, and exploration. For each combination of these four categories and the metal to be extracted, I compute the average area and assign it to the mine. In other words, a subsistence marble mine is weighted by the whole country's average area of all subsistence marble mines. Finally, Panel C allocates national exports of mining goods by the number of mines a township has of that good. In that specification, jade exposure is also allocated by the number of jade mines in a township, while it is done by area for all other specifications.

When choosing how to conduct this disaggregation, there is a trade-off between using all information available (area at face value) and having a more robust specification to outliers and potential misreporting (projected area, number of mines). I choose the main specification as a

[11] As for metals, most jade and gems exports from Myanmar go to China. Over the study period, China accounted for over 93% of the corresponding export value. I use exports to China to construct the control variable for export exposure in jade and gems. Unfortunately, the export data for opium are not broken down by destination. I, therefore, use estimates for opium exports to all destinations to construct the control for export exposure in opium (see Appendix B).

middle ground between these two options. Table A3 shows that the estimation results are fairly robust to this crucial modeling choice. The coefficients of interest are of similar magnitude in all specifications, showing a positive relationship between export exposure in metals and conflict events, which is statistically significant in almost all specifications.

Set of countries for the IV

For the instrumental variable strategy, I construct a measure of counterfactual export exposure using Chinese imports from low and middle-income countries (LMIC)[12] instead of Chinese imports from Myanmar (see Section 3).[13] The set of countries used to construct the instrument mirrors the one chosen by Autor et al. (2013), who use Chinese exports to other high-income countries. However, using another set of countries for the instrument is imaginable. Table A4 recomputes the instrumental variable regressions using three different sets of countries. The first two columns use imports from all other countries, including high-income countries. Columns (3) and (4) only use neighboring countries of China for the instrument. Finally, columns (5) and (6) consider all LMIC that are not neighbors of China. While some coefficients fail to reach statistical significance, they are all similar in size. If anything, using other sets of trading partners for the IV yields somewhat larger point estimates than the main specification..

Different conflict measures

This paper focuses on the number of violent conflict events in a township within a given year. The main specifications include all events in ACLED, except for those categorized as peaceful protests and strategic developments. Table A5 considers alternative conflict measures. Panel A is concerned with fatalities, while Panel B looks at fatal events, which means all events resulting in at least one fatality. Generally speaking, the direction of the effect is confirmed in all specifications with the baseline OLS effects being somewhat smaller, but statistically significant in all but one regression. However, the IV estimates are mostly insignificant and substantially smaller, compared with the main specification. This difference suggests that mining triggers violent conflict, but not always of the most deadly type. This is in line with the findings of Martin et al. (2008). They provide the intuition that increased international trade can substitute internal trade and therefore reduce the cost of conflict, but only if the scale of the conflict is limited such that trade can continue. Panel C considers the extensive margin of conflict. In this case, the dependent variable is a binary indicator of whether at least one reported conflict event exists for a township-year pair. A linear probability model confirms the main findings above: Export exposure is positively associated with conflict, especially in townships inhabited by ethnic minorities. Moreover, the IV specifications show larger coefficients, consistent with the OLS results being biased downwards due to reverse causality.

Government participation

Finally, Table A6 focuses on differential effects by the type of conflict parties. Panel A only includes events with the government forces as one of the conflict parties, while Panel B looks at

[12] All sets of countries exclude Myanmar.

[13] As reconciled trade data from BACI are used, exports from Myanmar to China and Chinese imports from Myanmar are equivalent.

conflicts in which none of the parties belongs to the government forces. It is unclear which type of conflict drives the main effect more. The OLS estimates are higher and statistically significant for events without government participation, while the IV estimates report higher coefficients for events with government participation. Nevertheless, this table contains interesting insights: First, the interaction effect seems stronger for conflict involving government forces. For this type of conflict, townships at least partially inhabited by the ethnic majority experienced a negative correlation between export exposure in mining goods and conflict. One plausible explanation for this finding is that the rebels struggle to find support to challenge the government's repression in these townships and increasingly so when the export business is booming. Another could be that increased revenues from mining exports enable the government to take increasingly repressive measures in these townships, shifting the balance of power to a point where challenging it is very costly. Second, the effect of export exposure in opium becomes substantially negative and highly significant for conflicts without government participation. While these results have to be taken with a grain of salt (see Section 2.4), they are in line with findings in the literature. Dal Bó and Dal Bó (2011) argue that one would expect a negative effect of prices on conflict for labor-intensive goods as the opportunity cost channel prevails over the rapacity channel. Dube and Vargas (2013) confirm this finding for coffee in Colombia. Opium production also requires a decent amount of manual labor. However, unlike coffee, it comes with the issue of the illegality of the sector. Basic intuition suggests that this illegality will mainly influence government actions, thereby potentially lessening the pacifying effect of the opportunity cost channel. For conflict between non-government actors, on the other hand, we would expect the economic trade-off between the rapacity and the opportunity channel to play more freely.

4.3 Placebo test

The OLS results in Section 4.1 show a positive correlation between export exposure in mining goods and conflict events. The instrumental variable strategy used above is the first way to validate these findings. However, exports from Myanmar increased in many different sectors during the 2010s. For example, Tanaka (2020) documents the positive effects of an increase in apparel exports on working conditions in the respective firms around Yangoon and Mandalay, the two biggest cities in the center of the country. To explore whether the effect on conflict is indeed plausibly driven by metals mined in Kachin and Shan states and not just an artifact of the general trade liberalization, I construct a placebo test.

Each metal used in this study is replaced by another good unrelated to metals. In order to represent the importance of the export sector of Myanmar in general and for trade with China in particular, I choose "placebo" goods to match the value of exports to China in the last pre-study year of 2011 as close as possible. For example, I replace copper with natural rubber, iron with tropical wood, and marble with dried fish. Table A7 describes the complete mapping.

These placebo trade flows are then used as if they would represent the actual trade flows of the metal to which they are mapped. For instance, I disaggregate exports (and Chinese imports) in tropical wood to townships according to their share of the national iron mining area using the same procedure as for the non-placebo specifications. Then, for each township, I sum

up disaggregated placebo trade flows over all metals existing there, producing a measure of counterfactual export exposure.

While the placebo goods are similarly important for exports to China in 2011, their evolution over time differs from their non-placebo counterparts. For some pairs, the export value grew faster in the placebo good. For others this value grew slower and peak export values occurred in different years. In particular, conflict around the mines should not impact export value development as the placebo goods are generally not produced there.

Table 2 shows the results of regressing conflict events on export exposure in placebo goods. All estimates are statistically insignificant and close to zero. IV estimates are positive, but with large standard errors. Therefore, the null hypothesis that placebo export exposure does not affect conflict can not be rejected. Given the construction of the measure, this is the behavior we would have expected. The results in Section 4.1 do not seem to be driven by trends in nationwide trade with China that are unrelated to mining.

Table 2: Placebo test

Dependent Variable:	ln(conflict events)					
Model:	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV
ln(export exposure metals)	0.12	-0.02	0.11	-0.07	0.14	-0.05
	(0.11)	(0.19)	(0.11)	(0.19)	(0.12)	(0.19)
Bamar × ln(export exposure metals)					-0.08	-0.20
					(0.12)	(0.75)
ln(export exposure jade)			0.18	0.18	0.18	0.19
			(0.22)	(0.22)	(0.22)	(0.23)
ln(export exposure opium)			-0.35	-0.37	-0.33	-0.32
			(0.26)	(0.26)	(0.27)	(0.29)
Cragg-Donald Wald F statistic		107.31		106.56		7.14

Note: The table shows regressions of violent conflict events on export exposure in placebo goods. Each primary sector mining good is replaced by another good that has an export value in the last pre-study year of 2011 that matches the one of the mining good as close as possible. I obtain export exposure by disaggregating nationwide exports in the placebo goods to the townships where the mining goods are produced, using the townships' share of the national production area in the mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. The export exposure in jade and in opium is computed as in Table 1. Specifications five and six consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports in the placebo goods from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

4.4 Night lights

In Kachin and Shan, from 2012 to 2020, export exposure in mining goods increased conflict in the townships where the mines are located. This finding is in line with the rapacity channel

dominating the opportunity cost channel (while the military capacity channel can potentially explain the geographic differences in the magnitude of the effect). As a final exercise, I assess the opportunity cost channel. Following the literature, I use VIIRS night light data to approximate income levels on a geographic resolution that is more disaggregated than the available income data (see, for example, Henderson et al. 2012). Compared to the data used above, I can work at an even finer scale for the night light data, as these data are available at a resolution of 500 $\times$ 500 meters.[14]

To leverage this precision, I recompute export exposure in mining goods at the mine level. Again, I disaggregate national exports to an individual mine, using the mine area as a weight for the allocation.[15] Equivalently to the procedure above, (log-) night light values are then regressed on the (log-) export exposure, using pixels as the unit of observation. I focus on mines that are already producing output. For conflict, we can expect the sheer presence of precious materials to matter. To expect an effect on income, however, something needs to be sold (or at least exploited). Apart from year fixed effects, I include a fixed effect for every pixel in a panel regression. Standard errors are two-way clustered at the year and the pixel level.

I split the sample into seven parts, depending on the distance of a pixel to the closest mine: 0 to 1 km, 1 to 2 km, 2 to 3 km, 3 to 5 km, 5 to 10 km, 10 to 15km, and 15 to 20 km.[16] This allows me to assess whether the economic effects of higher export exposure only affect the mine and its immediate surroundings or if they also spread further into the regions that host the mines.

The red points in Figure 3 show the regression coefficients for all pixels in Kachin and Shan that are in the respective distance brackets. A statistically significant effect of export exposure on night light values can be found in the immediate neighborhood of the mines. For the first bracket, a 10 percent higher level of export exposure is associated with a 0.6 percent higher level of night lights. This effect might be due to the workers employed at the mines and local businesses directly involved.

On the other hand, spillover effects are very small. The effect still seems to be positive for pixels between 1 and 2 kilometers from the mines but fails to meet significance at the 10% level. Furthermore, the estimates are statistically insignificant for all distance brackets above 2 kilometers, with point estimates close to zero. The weak spillover effects are consistent with reports that argue that the gains from metal exploitation rarely benefit the local population. Instead, they are often transferred to more central parts of Myanmar, benefiting political leaders, army officials, or investors abroad. Global Witness (2015) provides a vivid example for the case of jade.

I redo the same analysis taking the subset of pixels that are in parts of Kachin and Shan that are inhabited by Bamar (blue points in Figure 3) and in parts that are inhabited exclusively by

[14] More precisely, cell size is 500 $\times$ 500 meters at the equator and somewhat larger in the case of Myanmar.

[15] I cap mine area at the area of the largest mine classified as producing, consistent with the main specification above.

[16] I measure the distance from the centroid of a pixel to the centroid of a mine. A pixel is exclusively assigned to the bracket that corresponds to the closest distance, even if other mines exist further away. In case multiple mines are situated within the same distance bracket, I sum up their respective export exposures.

Figure 3: Elasticity of night lights to export exposure in mining, by distance from the mine

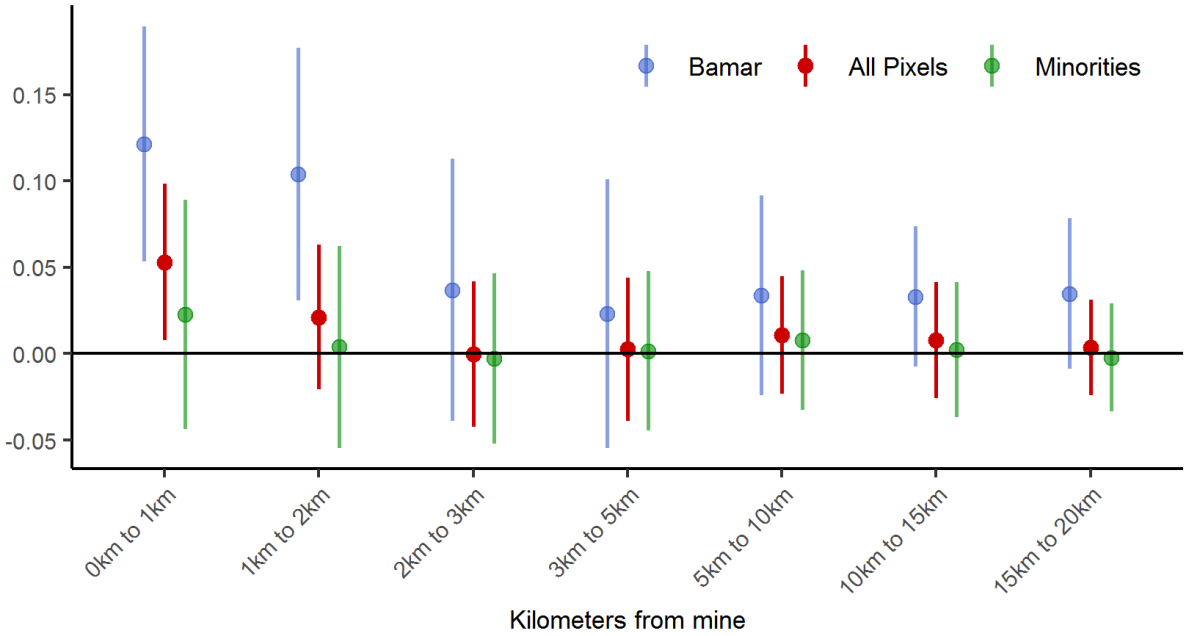

Note: Each point represents the estimated coefficient of a regression of night lights on export exposure in primary sector mining goods. The units of observation are pixels (500m × 500m). The regressions differ by the distance between a pixel and its closest mine (7 categories) and by whether all pixels in Kachin and Shan are considered (red) or only the subset in areas inhabited by Bamar (blue) or exclusively by ethnic minorities (green). The panels consist of between 7,416 (0km to 1km, Bamar) and 982,973 (15km to 20km, all pixels) observations. All specifications include pixel fixed effects and year fixed effects. I obtain export exposure by disaggregating nationwide exports to the mines where the goods are produced, using the mines' share of the national production area. Mines categorized as exploration are excluded. If two mines are in the same distance bracket, I sum up their export exposures. I only include pixels in the regressions that correspond to their closest bracket. I take logs for export exposure and night light to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The lines represent confidence intervals at the 90% level, constructed from standard errors, which are two-way clustered by mine and by year.

minorities (green points). Minority areas seem to profit less from export exposure in mining than areas with a presence of the ethnic majority. This finding affects both the direct surroundings of the mines and the areas in the regions around them. For the pixels in Bamar areas, the elasticity is large and statistically significant for the first two brackets and higher than the main specification for all brackets. The pixels in territories exclusively inhabited by minorities show no significant effect in any bracket, with a point estimate of essentially zero already from the second bracket.

5 Conclusion

I use a two-way panel data model on a township-year basis to estimate the effect of trade in natural resources on conflict. While the literature tends to find that trade attenuates conflict, various studies find that higher world-market prices for mining goods increase conflict. This group of goods is, therefore, an obvious candidate to test the limits of the pacifying effect of trade. Using novel data on mining licenses, I analyze the effect of a trade shock in the case of

Myanmar. Most of the trade in mining goods happens with the neighboring country of China. While mining companies in Myanmar depend on Chinese demand, Myanmar is only a minor trading partner of China.

As an identification strategy, I first disaggregate nationwide export values for different metals to the townships where these metals are exploited. Then, as a second step, I construct an instrument using trade flows of the same metals from other low and middle-income countries to China. This procedure essentially mirrors the identification strategy of Autor et al. (2013), who use Chinese exports to high-income countries other than the United States to construct their instrument.

I focus on Kachin and Shan, the two states within Myanmar that share a border with China. Both states are conflict-ridden, including many incidents between Myanmar's government and rebel groups; often made up of ethnic minorities. Both states are crucial for the mining sector of Myanmar. Moreover, they produce a considerable amount of opium and jade; two factors for which I control.

The results suggest that trade in primary sector mining goods is indeed associated with conflict. The elasticity between violent conflict events and export exposure in metals is about 0.46 for the OLS specification and about 0.56 for the IV specification that controls for reverse causality. There is heterogeneity in the results concerning the presence of the ethnic majority. The positive relation appears to be predominantly driven by townships where the ethnic majority is absent.

A possible explanation for this finding could be a higher level of military control over areas inhabited by the ethnic majority. Crost and Felter (2019) consider the case of bananas in the Philippines and find that increased export exposure is more likely to trigger conflict in townships that are not fully controlled by one conflict party. Arezki and Brueckner (2021) show that the positive relation between export exposure and conflict disappears for countries with high military expenditures. According to the World Development Indicators from the World Bank, Myanmar is 17th of 140 countries when it comes to military expenditure as a percentage of GDP,[17] so one would expect their findings to be relevant for this setting. Only at the fringes of the country might the military be weak enough to be seriously challenged by other actors.

Violent conflict has severe costs. Nevertheless, there is, of course, also an upside to an increased trade volume, at least in theory. First, it should create jobs, allowing people to become employed in lawful sectors. Second, profiting from the sale of natural resources via taxes or direct involvement would allow the government to inject funds into the public treasury that could be used for the benefit of the people. Unfortunately, only a tiny part of Myanmar's revenue from the sale of natural resources seems to be allocated to the local population in the regions where these resources are mined (Global Witness, 2015). This lack of local benefit can be validated using night light data as a proxy for economic activity. An increase in a mine's export exposure is indeed correlated with an increase in night light activity. However, this increase is very local

[17] Military spending was on average 3.5% of GDP over the years 2012-2020. For this statistic, I only considered countries with non-missing values for all 9 years. Data downloaded 28.09.2022 from `https://databank.worldbank.org/source/world-development-indicators`.

around the mines. The data cannot confirm spillover effects on the surrounding areas. Moreover, the increase in night lights is smaller, and declines even quicker, in minority areas. In this regard, too, the recent developments around the military coup of February 2021 do not give rise to optimism.

6 Bibliography

Aragón, F.M. and Rud, J.P., 2016. Polluting industries and agricultural productivity: Evidence from mining in Ghana, *The Economic Journal*, 126 (597), 1980–2011.

Arezki, R. and Brueckner, M., 2021. Natural resources and civil conflict: The role of military expenditures, *Journal of Risk and Financial Management*, 14, 575.

Autor, D., Dorn, D., Hanson, G., and Majlesi, K., 2020. Importing political polarization? The electoral consequences of rising trade exposure, *American Economic Review*, 110 (10), 3139–3183.

Autor, D.H., Dorn, D., and Hanson, G.H., 2013. The China syndrome: Local labor market effects of import competition in the United States, *American Economic Review*, 103 (6), 2121–2168.

Bartik, T.J., 1991. *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research.

Berman, N., Couttenier, M., Rohner, D., and Thoenig, M., 2017. This mine is mine! How minerals fuel conflicts in Africa, *American Economic Review*, 107 (6), 1564–1610.

Brülhart, M., Cadot, O., and Himbert, A., 2019. Let there be light: Trade and the development of border regions, Discussion Paper, C.E.P.R.

Calì, M. and Miaari, S.H., 2015. Trade, employment and conflict: Evidence from the second intifada, Report, Overseas Development Institute.

Candau, F., Gbandi, T., and Guepie, G., 2022. Beyond the income effect of international trade on ethnic wars in Africa, *Economics of Transition and Institutional Change*, 30 (3), 517–534.

Chen, X. and Nordhaus, W., 2015. A test of the new VIIRS lights data set: Population and economic output in Africa, *Remote Sensing*, 7 (4), 4937–4947.

Colantone, I. and Stanig, P., 2018. Global competition and Brexit, *American Political Science Review*, 112 (2), 201–218.

Crost, B. and Felter, J.H., 2019. Export crops and civil conflict, *Journal of the European Economic Association*, 18 (3), 1484–1520.

Crost, B. and Felter, J.H., 2020. Extractive resource policy and civil conflict: Evidence from mining reform in the Philippines, *Journal of Development Economics*, 144.

Dal Bó, E. and Dal Bó, P., 2011. Workers, warriors, and criminals: Social conflict in general equilibrium, *Journal of the European Economic Association*, 9 (4), 646–677.

Dauth, W., Findeisen, S., and Suedekum, J., 2017. Trade and manufacturing jobs in Germany, *American Economic Review*, 107 (5), 337–342.

Dube, O. and Vargas, J.F., 2013. Commodity price shocks and civil conflict: Evidence from Colombia, *Review of Economic Studies*, 80 (4), 1384–1421.

Elvidge, C., Baugh, K., Zhizhin, M., and Hsu, F.C., 2013. Why VIIRS data are superior to DMSP for mapping nighttime lights, *Proceedings of the Asia-Pacific Advanced Network*, 35, 62–69.

Elvidge, C.D., Zhizhin, M., Ghosh, T., Hsu, F.C., and Taneja, J., 2021. Annual time series of global VIIRS nighttime lights derived from monthly averages: 2012 to 2019, *Remote Sensing*, 13 (5), 922.

Ferraz, E., Soares, R., and Vargas, J., 2021. Unbundling the relationship between economic shocks and crime, Working Paper, IZA Institute of Labor Economics.

Giménez-Gómez, J.M. and Zergawu, Y.Z., 2018. The impact of social heterogeneity and commodity price shocks on civil conflicts, *Journal of Policy Modeling*, 40 (5), 959–997.

Glick, R. and Taylor, A.M., 2010. Collateral damage: Trade disruption and the economic impact of war, *The Review of Economics and Statistics*, 92 (1), 102–127.

Global Witness, 2015. Jade: Myanmar's "big state secret", Report, Global Witness, `https://www.globalwitness.org/en/campaigns/oil-gas-and-mining/myanmarjade/` (last accessed: 20.09.2022.

Henderson, J.V., Storeygard, A., and Weil, D.N., 2012. Measuring economic growth from outer space, *American Economic Review*, 102 (2), 994–1028.

Janus, T. and Riera-Crichton, D., 2015. Economic shocks, civil war and ethnicity, *Journal of Development Economics*, 115, 32–44.

Kramer, T., 2017. The current state of counternarcotics policy and drug reform debates in Myanmar, *Journal of Drug Policy Analysis*, 10 (1), 721–800.

Martin, P., Thoenig, M., and Mayer, T., 2008. Civil wars and international trade, *Journal of the European Economic Association*, 6 (2), 541–550.

McManus, T.C. and Schaur, G., 2016. The effects of import competition on worker health, *Journal of International Economics*, 102, 160–172.

MCRB, 2022. MEITI licenses explorer for 6th MEITI summary data report, Online, Myanmar Centre for Responsible Business, `https://meiti-map.org` (last accessed: 07.03.2023).

MEITI, 2015. Myanmar Extractive Industries Transparency Initiative report for the period April 2013 - March 2014: Oil, gas and mining sectors, Report, Moore Stephens LLP, `https://eiti.org/documents/myanmar-2013-2014-eiti-report` (last accessed: 20.09.2022).

MEITI, 2020. 6th MEITI summary data report, Report, The Republic of the Union of Myanmar, `https://myanmareiti.org/en/publication/6th-meiti-summary-data-report` (last accessed: 07.03.2023).

Morelli, M. and Rohner, D., 2015. Resource concentration and civil wars, *Journal of Development Economics*, 117, 32–47.

Oak, Y.N., 2018. Even with new data, valuing Myanmar's jade industry remains a challenge, https://openjadedata.org/stories/how_much_jade_worth.html (last accessed: 12.10.2022).

Qureshi, M.S., 2013. Trade and thy neighbor's war, *Journal of Development Economics*, 105, 178–195.

Raleigh, C., Linke, A., Hegre, H., and Karlsen, J., 2010. Introducing ACLED: An armed conflict location and event dataset: Special data feature, *Journal of Peace Research*, 47 (5), 651–660.

Tanaka, M., 2020. Exporting sweatshops? Evidence from Myanmar, *The Review of Economics and Statistics*, 102 (3), 442–456.

UNODC, 2011. South-East Asia opium survey 2011 Lao PDR, Myanmar, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/southeastasiaandpacific/2011/12/ops-2011/Opium_Survey_2011_-_Full.pdf (last accessed: 27.09.2022).

UNODC, 2012. South-East Asia opium survey 2012 Lao PDR, Myanmar, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/crop-monitoring/sea/SouthEastAsia_Report_2012_low.pdf (last accessed: 27.09.2022).

UNODC, 2013. Southeast Asia opium survey 2013 Lao PDR, Myanmar, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/southeastasiaandpacific/Publications/2013/SEA_Opium_Survey_2013_web.pdf (last accessed: 27.09.2022).

UNODC, 2014. Southeast Asia opium survey 2014 Lao PDR, Myanmar, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/southeastasiaandpacific/Publications/2014/ops/SE_ASIA_opium_poppy_2014_em.pdf (last accessed: 27.09.2022).

UNODC, 2015. Southeast Asia opium survey 2014 Lao PDR, Myanmar, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/crop-monitoring/sea/Southeast_Asia_Opium_Survey_2015_web.pdf (last accessed: 27.09.2022).

UNODC, 2017a. Evidence for enhancing resilience to opium poppy cultivation in Shan state, Myanmar. implications for alternative development, peace, and stability, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/crop-monitoring/sea/2016_Myanmar_Shan_Opium_Poppy_web.pdf (last accessed: 27.09.2022).

UNODC, 2017b. Myanmar opium survey, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/southeastasiaandpacific/Publications/2017/Myanmar_Opium_Survey_2017_web.pdf (last accessed: 27.09.2022).

UNODC, 2018. Myanmar opium survey 2018. cultivation, production and implications, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/

documents/crop-monitoring/Myanmar/Myanmar_Opium_Survey_2018-web.pdf (last accessed: 27.09.2022).

UNODC, 2019. Myanmar opium survey 2019. cultivation, production and implications, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/southeastasiaandpacific/Publications/2020/Myanmar_Opium_Survey_2019.pdf (last accessed: 27.09.2022).

UNODC, 2021. Myanmar opium survey 2020. cultivation, production and implications, Report, United Nations Office on Drugs and Crime, https://www.unodc.org/documents/crop-monitoring/Myanmar/Myanmar_Opium_survey_2020.pdf (last accessed: 27.09.2022).

Vogt, M., Bormann, N.C., Rüegger, S., Cederman, L.E., Hunziker, P., and Girardin, L., 2015. Integrating data on ethnicity, geography, and conflict: The ethnic power relations data set family, *Journal of Conflict Resolution*, 59 (7), 1327–1342.

Woods, K.M., 2018. *The Conflict Resource Economy and Pathways to Peace in Burma*, United States Institute of Peace Washington, DC.

A Additional descriptive statistics

Figure A1: Number of conflict fatalities between 2012 and 2020

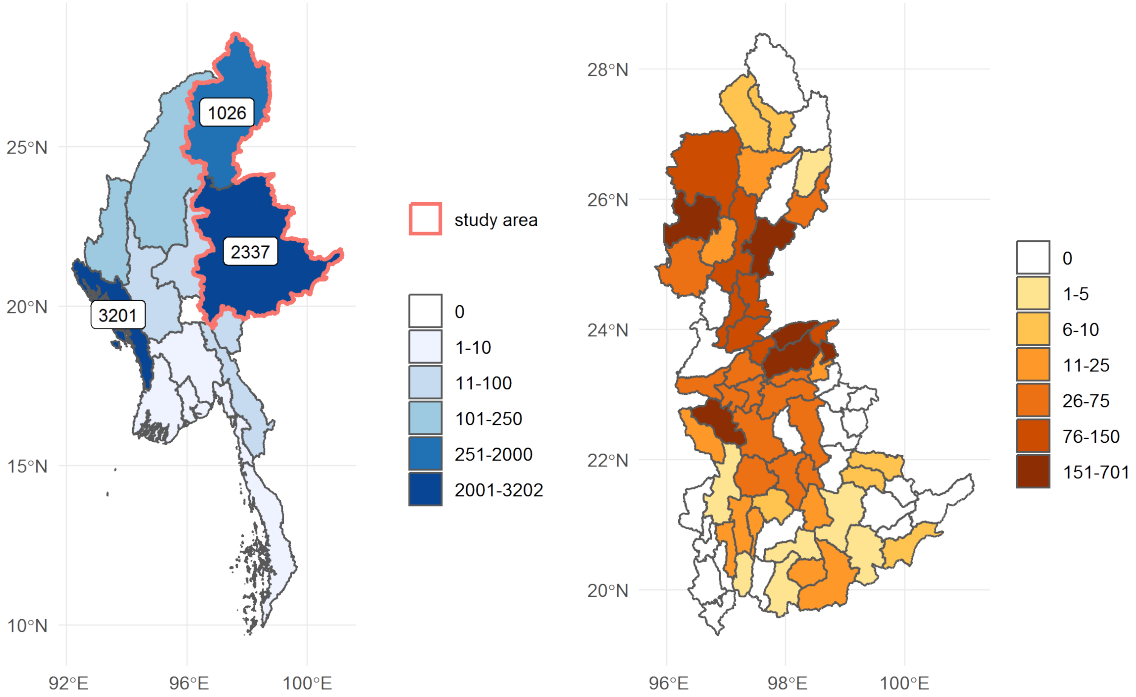

(a) by state

(b) by township in Kachin and Shan states

Figure A2: Presence of the ethnic majority (Bamar)

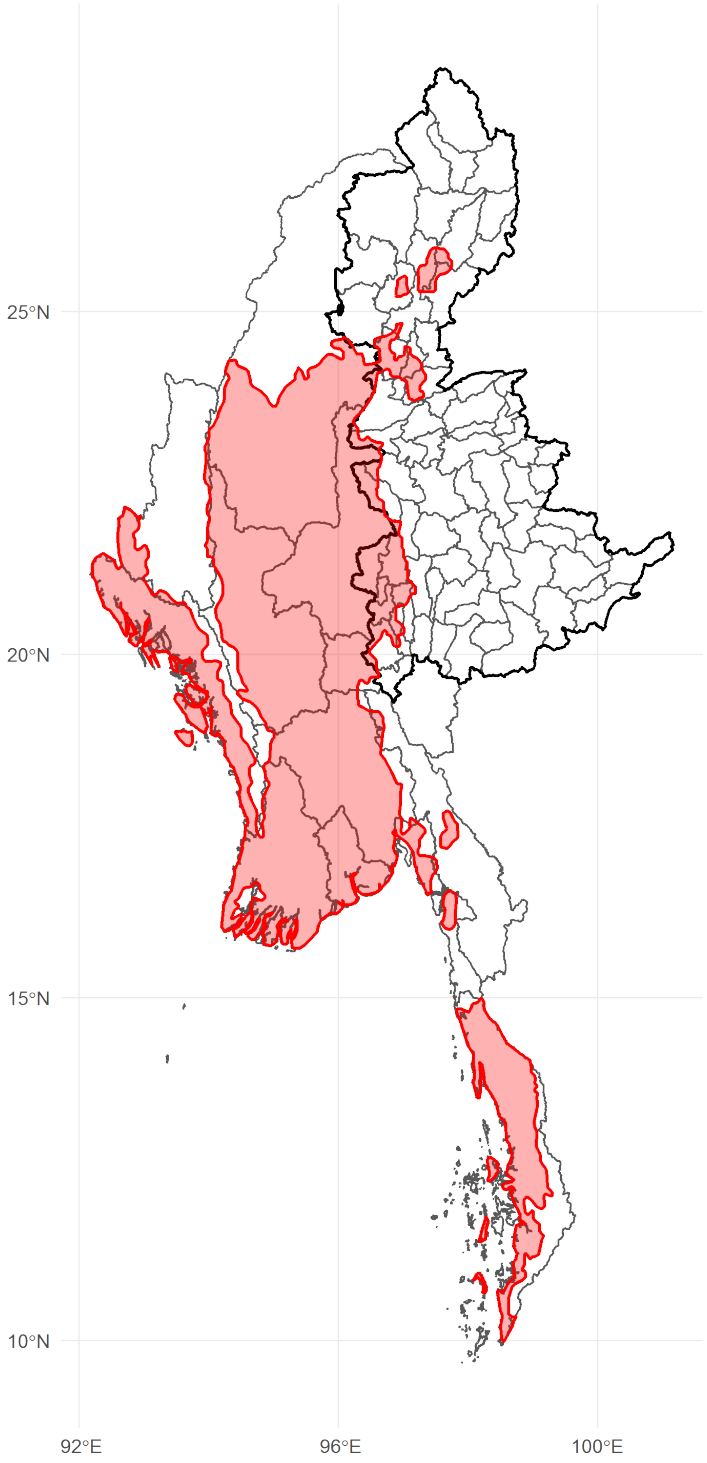

Table A1: Mapping harmonized system (96) codes to metals.

HS6 Code	Description	Assigned to
261710	Antimony ores and concentrates	Antimony
251110	Barium sulphate (barytes): natural	Baryte
261000	Chromium ores and concentrates	Chromium
270112	Coal: bituminous, whether or not pulverised, but not agglomerated	Coal
270119	Coal: (other than anthracite and bituminous), whether or not pulverised but not agglomerated	Coal
260300	Copper ores and concentrates	Copper
262030	Ash and residues: (not from the manufacture of iron or steel), containing mainly copper	Copper
740400	Copper: waste and scrap	Copper
251810	Dolomite: (not calcined), roughly trimmed or merely cut, by sawing or otherwise, into blocks or slabs of a rectangular (including square) shape	Dolomite
252010	Gypsum: anhydrite	Gypsum
260111	Iron ores and concentrates: non-agglomerated	Iron
260112	Iron ores and concentrates: agglomerated (excluding roasted iron pyrites)	Iron
260120	Iron pyrites: roasted	Iron
261900	Slag, dross: (other than granulated slag), scalings and other waste from the manufacture of iron or steel	Iron
720429	Ferrous waste and scrap: of alloy steel (excluding stainless)	Iron
720449	Ferrous waste and scrap: n.e.s. in heading no. 7204	Iron
260700	Lead ores and concentrates	Lead & Zinc
260800	Zinc ores and concentrates	Lead & Zinc
262019	Ash and residues: (not from the manufacture of iron or steel), containing mainly zinc, other than hard zinc spelter	Lead & Zinc
262020	Ash and residues: (not from the manufacture of iron or steel), containing mainly lead	Lead & Zinc
780200	Lead: waste and scrap	Lead & Zinc
252100	Limestone flux: limestone and other calcareous stone, of a kind used for the manufacture of lime or cement	Limestone
260200	Manganese ores and concentrates, including ferruginous manganese ores and concentrates with a manganese content of 20% or more, calculated on the dry weight	Manganese Dioxide
251511	Marble and travertine: having a specific gravity of 2.5 or more, crude or roughly trimmed by sawing or otherwise, into blocks or slabs of a rectangular (including square) shape	Marble
251512	Marble and travertine: merely cut, by sawing or otherwise, into blocks or slabs of a rectangular (including square) shape, having a specific gravity of 2.5 or more	Marble
251741	Stones: of marble, in granules, chippings and powder, whether or not heat-treated	Marble Marble
250510	Sands: natural, silica and quartz sands, whether or not coloured	Quartzite
250610	Quartz: other than natural sands	Quartzite
250621	Quartzite: crude or roughly trimmed	Quartzite

250629	Quartzite: cut, by sawing or otherwise, into blocks or slabs of a rectangular (including square) shape, (excluding crude or roughly trimmed)	Quartzite
260900	Tin ores and concentrates	Tin & Tungsten
261100	Tungsten ores and concentrates	Tin & Tungsten
800200	Tin: waste and scrap	Tin & Tungsten

B Approximating export exposure in opium

The Myanmar Opium Survey 2020 (UNODC, 2021, p. 3) includes a map which describes the average cultivation density of opium between 2014 and 2020. This map contains four classes of density. To get a measure of opium production by township, I trace this map and take the following steps:

1. For each class, the UNODC report gives a bracket of cultivation density. For the low (0 - 0.01 ha/km^2), medium (0.01 - 0.1 ha/km^2), and high (0.1 - 1 ha/km^2) class, I take the average of this bracket (0.005, 0.055, and 0.55 ha/km^2 respectively). For the very high class ($>$ 1 ha/km^2), I assume a value of 2 ha/km^2.

2. I assume that areas outside the survey area of the UNODC report have no opium production. UNODC (2021, p. 33) provides support for this assumption.

3. I compute the overlapping area between each township and each of the four density classes.

4. I estimate the total opium cultivation in a township by multiplying the area of the township that is covered by a class with that class's cultivation density and then summing up the corresponding values for all classes that are present in the township.

5. To calculate the share of opium production of a township, I divide that township's opium production by the sum of opium production across all townships (divided by 31900/33100, as UNODC (2021) estimates that 1'200 ha of the opium cultivation area located outside Kachin and Shan states).

I then use this share of opium production by township to disaggregate the value generated from each year's nationwide export of opium. Estimates of the latter can also be found in the yearly UNODC reports. Unfortunately, the reports for the years 2016 (UNODC, 2017a) and 2017 UNODC (2017b) do not contain information about the value of opium production. Moreover, not all reports estimate the same matter. To make the different numbers comparable, I take the following additional assumptions:

1. The reports concerning the years 2012-2015 (UNODC (2012), UNODC (2013), UNODC (2014), UNODC (2015)) always provide point estimates as well as brackets of values for the relevant variables. This is not true for the years 2018-2020 (UNODC (2018), UNODC (2019), UNODC (2021)), which only report brackets for some values. In these cases, I take the middle of the brackets as point estimates. I also use the middle of the bracket to correct the farm-gate value in 2020. In this case the report provides a mid-point which is, however, outside of the bracket.

2. To make the values for different years comparable, I convert them to the farm-gate value in a first step. While the reports contain farm-gate values for the years 2018-2020, they only have an estimate of the "total potential wholesale value" for 2012-2015. However, there are estimates for both terms in 2011, once from the 2011 report (UNODC, 2011, farm-gate value) and once from the 2012 report (wholesale value). As the latter is 1.2 times as large as the former, I divide the wholesale value by 1.2 to get estimates for the

farm-gate value for 2012-2015.

3. The reports for the years 2018-2020 include an estimate of the "farm-gate value of opium" as well an estimate of the "value of opiates potentially available for export" (at the border). On average, the latter is ≈ 11.3 times as large as the former. As the interest of this exercise is export exposure, I scale the farm-gate values for 2012-2015 by ≈ 11.3.

4. Finally, I linearly interpolate the values of 2015 and 2018 to get estimates for the missing years 2016 and 2017.

C Robustness checks

Table A2: Control variables

Dependent Variable:	ln(conflict events)			
Model:	(1) OLS	(2) IV	(3) OLS	(4) IV
Panel A: Only townships without jade or gems				
ln(export exposure metals)	0.42**	0.47*	0.67**	0.97***
	(0.17)	(0.22)	(0.22)	(0.28)
Bamar × ln(export exposure metals)			-0.72**	-1.49**
			(0.28)	(0.60)
ln(export exposure opium)	-0.46	-0.47	-0.33	-0.20
	(0.27)	(0.27)	(0.29)	(0.33)
Cragg-Donald Wald F statistic		192.81		93.51
Panel B: Not controlling for opium export exposure				
ln(export exposure metals)	0.42**	0.51*	0.74**	1.06***
	(0.17)	(0.26)	(0.22)	(0.28)
Bamar × ln(export exposure metals)			-0.86**	-1.57**
			(0.26)	(0.55)
ln(export exposure jade)	0.15	0.15	0.14	0.12
	(0.21)	(0.21)	(0.20)	(0.20)
Cragg-Donald Wald F statistic		206.78		98.50

Note: Panel A excludes three townships where MEITI reports licenses related to jade and/or gemstones. Panel B includes all townships, but does not control for export exposure in opium. The table shows regressions of violent conflict events on export exposure in primary sector mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. I obtain export exposure by disaggregating nationwide exports to the townships where the goods are produced, using the townships' share of the national production area. The export exposure in jade and in opium is computed in the same way. Specifications three and four consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A3: Allocation of national exports to townships

Dependent Variable:	ln(conflict events)			
Model:	(1) OLS	(2) IV	(3) OLS	(4) IV
Panel A: Exports allocated by mining area				
ln(export exposure metals)	0.47**	0.48	0.76**	1.02**
	(0.18)	(0.34)	(0.24)	(0.38)
Bamar × ln(export exposure metals)			-0.80**	-1.55**
			(0.31)	(0.65)
ln(export exposure jade)	0.16	0.16	0.14	0.13
	(0.20)	(0.21)	(0.20)	(0.20)
ln(export exposure opium)	-0.40	-0.40	-0.29	-0.19
	(0.26)	(0.26)	(0.28)	(0.30)
Cragg-Donald Wald F statistic		163.28		75.15
Panel B: Exports allocated by projected mining area				
ln(export exposure metals)	0.39*	0.56**	0.55**	1.08***
	(0.18)	(0.23)	(0.21)	(0.24)
Bamar × ln(export exposure metals)			-0.49	-1.67**
			(0.30)	(0.60)
ln(export exposure jade)	0.15	0.14	0.15	0.12
	(0.20)	(0.21)	(0.20)	(0.20)
ln(export exposure opium)	-0.40	-0.42	-0.33	-0.17
	(0.26)	(0.26)	(0.27)	(0.31)
Cragg-Donald Wald F statistic		183.14		87.72
Panel C: Exports allocated by number of mines				
ln(export exposure metals)	0.32*	0.48*	0.52**	0.93***
	(0.16)	(0.21)	(0.20)	(0.26)
Bamar × ln(export exposure metals)			-0.57*	-1.41**
			(0.25)	(0.56)
ln(export exposure jade)	0.16	0.15	0.15	0.12
	(0.20)	(0.20)	(0.20)	(0.19)
ln(export exposure opium)	-0.42	-0.45	-0.34	-0.24
	(0.26)	(0.26)	(0.27)	(0.31)
Cragg-Donald Wald F statistic		187.92		91.66

Note: Panel A disaggregates nationwide exports by mining area at face value, without capping outliers beyond the largest mine classified as producing. Panel B assigns to each mine the area of an average mine of the same type and metal. Panel C uses the number of mines to disaggregate exports. Export exposure in jade is disaggregated by area in Panels A and B, and by the number of mines in Panel C. The table shows regressions of violent conflict events on export exposure in primary sector mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. I obtain export exposure by disaggregating nationwide exports to the townships where the goods are produced, using the townships' share of the national production area. The export exposure in jade and in opium is computed in the same way. Specifications five and six consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A4: Different sets of countries for the instrument

Dependent Variable:	ln(conflict events)					
Subset:	All countries		Neighbors		Lmic ex neigh.	
Model:	(1) IV	(2) IV	(3) IV	(4) IV	(5) IV	(6) IV
ln(export exposure metals)	0.64	1.16**	0.67	1.18*	0.69**	1.11**
	(0.35)	(0.36)	(0.45)	(0.60)	(0.29)	(0.34)
Bamar × ln(export exposure metals)		-1.48**		-1.43		-1.32**
		(0.59)		(0.83)		(0.55)
ln(export exposure jade)	0.14	0.11	0.13	0.11	0.13	0.12
	(0.21)	(0.20)	(0.20)	(0.18)	(0.20)	(0.20)
ln(export exposure opium)	-0.46	-0.21	-0.46	-0.23	-0.46	-0.24
	(0.26)	(0.31)	(0.25)	(0.31)	(0.27)	(0.32)
Cragg-Donald Wald F statistic	164.48	76.46	169.36	83.28	178.42	88.45

Note: Columns (1) and (2) use Chinese imports from all countries (other than Myanmar) for the construction of the instrument. Columns (3) and (4) use imports from the neighboring countries of China. Columns (5) and (6) use imports from low and middle-income countries that are not neighbors of China. The table shows regressions of violent conflict events on export exposure in primary sector mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. I obtain export exposure by disaggregating nationwide exports to the townships where the goods are produced, using the townships' share of the national production area. The export exposure in jade and in opium is computed in the same way. Specifications two, four, and six consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A5: Conflict measures

Model:	(1) OLS	(2) IV	(3) OLS	(4) IV
Dependent Variable:	ln(fatalities)			
Panel A: Fatalities				
ln(export exposure metals)	0.41**	0.32	0.51*	0.55
	(0.17)	(0.40)	(0.23)	(0.38)
Bamar × ln(export exposure metals)			-0.28	-0.70
			(0.38)	(0.74)
ln(export exposure jade)	0.15	0.16	0.15	0.15
	(0.26)	(0.28)	(0.26)	(0.27)
ln(export exposure opium)	-0.30	-0.29	-0.26	-0.17
	(0.24)	(0.23)	(0.21)	(0.22)
Cragg-Donald Wald F statistic		203.16		98.49
Dependent Variable:	ln(fatal events)			
Panel B: Fatal events				
ln(export exposure metals)	0.33**	0.30	0.40	0.50*
	(0.14)	(0.27)	(0.21)	(0.22)
Bamar × ln(export exposure metals)			-0.19	-0.59
			(0.34)	(0.56)
ln(export exposure jade)	0.10	0.10	0.09	0.09
	(0.22)	(0.23)	(0.22)	(0.23)
ln(export exposure opium)	-0.37*	-0.36*	-0.34*	-0.26
	(0.17)	(0.17)	(0.16)	(0.17)
Cragg-Donald Wald F statistic		203.16		98.49
Dependent Variable:	has conflict event			
Panel C: Binary indicator, at least one violent conflict event				
ln(export exposure metals)	0.10*	0.16*	0.14**	0.27**
	(0.05)	(0.09)	(0.06)	(0.11)
Bamar × ln(export exposure metals)			-0.11*	-0.32*
			(0.05)	(0.17)
ln(export exposure jade)	0.05	0.04	0.04	0.04
	(0.03)	(0.03)	(0.03)	(0.03)
ln(export exposure opium)	-0.04	-0.05	-0.02	0.00
	(0.07)	(0.07)	(0.08)	(0.10)
Cragg-Donald Wald F statistic		203.16		98.49

Note: Panel A uses the number of fatalities in a township and year as the dependent variable. Panel B uses the number of events with at least one fatality. Panel C uses a binary variable, indicating whether there has been at least one violent conflict event. The main variable of interest is export exposure in primary sector mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. I obtain export exposure by disaggregating nationwide exports to the townships where the goods are produced, using the townships' share of the national production area. The export exposure in jade and in opium is computed in the same way. Specifications three and four consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * p $<$ 0.10 ** p $<$ 0.05 *** p $<$ 0.01.

Table A6: Government participation

Dependent Variable:	ln(conflict events)			
Model:	(1) OLS	(2) IV	(3) OLS	(4) IV
Panel A: Events with government participation				
ln(export exposure metals)	0.19	0.30	0.42	0.83
	(0.19)	(0.43)	(0.29)	(0.52)
Bamar × ln(export exposure metals)			-0.63	-1.59**
			(0.35)	(0.58)
ln(export exposure jade)	0.17	0.16	0.16	0.14
	(0.20)	(0.21)	(0.20)	(0.21)
ln(export exposure opium)	-0.28	-0.29	-0.17	-0.03
	(0.29)	(0.30)	(0.31)	(0.34)
Cragg-Donald		203.16		98.49
Panel B: Events without government participation				
ln(export exposure metals)	0.69***	0.51*	0.93**	0.74*
	(0.20)	(0.27)	(0.28)	(0.39)
Bamar × ln(export exposure metals)			-0.67	-0.69
			(0.36)	(0.50)
ln(export exposure jade)	-0.15***	-0.14**	-0.16***	-0.14***
	(0.03)	(0.04)	(0.02)	(0.03)
ln(export exposure opium)	-0.78***	-0.75***	-0.67***	-0.64***
	(0.19)	(0.18)	(0.17)	(0.14)
Cragg-Donald		203.16		98.49

Note: Panel A includes all violent events in which at least one of the conflict parties was coded as state forces by ACLED, including military and police forces. Panel B includes all other violent events, where none of the parties was coded as belonging to the state forces. The table shows regressions of violent conflict events on export exposure in primary sector mining goods. The balanced panel consists of 73 townships in Kachin and Shan and 9 years (657 observations). All specifications include township fixed effects and year fixed effects. The parentheses show standard errors, which are two-way clustered by township and by year. I obtain export exposure by disaggregating nationwide exports to the townships where the goods are produced, using the townships' share of the national production area. The export exposure in jade and in opium is computed in the same way. Specifications three and four consider heterogeneous effects by including an interaction term, with Bamar being an indicator variable for the presence of the ethnic majority in a township. For the IV specifications, I construct a shift-share instrument using Chinese imports from other low and middle-income countries, multiplied by the same shares of national production area as above. I take logs for all export exposures and conflict events to interpret the coefficients as elasticities. To keep observations with zero values, I add 0.1 to all variables before taking logs. The levels of significance are * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

Table A7: Mapping of metals to goods of similar value for placebo test

Metal	HS6 code	and description of matched placebo good	Mio USD (2011*) Metal	Placebo
Antimony	080300	Fruit, edible: bananas, (including plantains), fresh or dried	7.40	7.57
Baryte	620291	Anoraks (including ski-jackets), wind-cheaters, wind-jackets and similar articles: women's or girls', of wool or fine animal hair, other than those of heading no. 6204 (not knitted or crocheted)	0.04	0.04
Chromium	530310	Jute and other textile bast fibres: raw or retted, but not spun, (excluding flax, hemp (cannabis sativa L.), and ramie)	0.02	0.01
Coal	240210	Cigars, cheroots and cigarillos: containing tobacco including the weight of every band, wrapper or attachment thereto	0.01	0.01
Copper	400129	Rubber: natural (excluding latex, technically specified natural rubber and smoked sheets), in primary forms or in plates, sheets or strip	10.53	10.54
Dolomite	030749	Molluscs: cuttle fish and squid, frozen, dried, salted or in brine (whether in shell or not)	0.42	0.42
Gypsum	841990	Machinery, plant and laboratory equipment: parts of equipment for treating materials by a process involving a change of temperature	0.00	0.00
Iron	440349	Wood, tropical: (as specified in subheading note 1, chapter 44, customs tariff), n.e.s. in item no. 4403.41, in the rough, whether or not stripped of bark or sapwood, or roughly squared, untreated	87.47	87.94
Lead & Zinc	121299	Vegetable products (including unroasted chicory roots, chicorium intybus sativum variety): n.e.s. in chapter 12, fresh, chilled, frozen or dried, ground or unground, primarily for human consumption	21.48	21.92
Limestone	070490	Vegetables, brassica: edible, n.e.s. in heading no. 0704, fresh or chilled	0.00	0.00
Manganese Dioxide	400122	Rubber: technically specified natural rubber (TSNR), in primary forms or in plates, sheets or strip (excluding latex and smoked sheets)	43.57	39.82
Marble	030559	Fish: dried (whether or not salted but not smoked), n.e.s. in item no. 0305.51	1.10	1.13
Quartzite	410619	Leather: goat or kid skin, without hair on, tanned or retanned but not further prepared, whether or not split, n.e.s. in heading no. 4106 (not pre-tanned), excluding leather of heading no. 4108 and 4109	0.38	0.39
Tin & Tungsten	100510	Cereals: maize (corn), seed	31.82	35.60

Note: Matched such that the difference in export value to China in the last pre-study year of 2011 is minimized. The following two goods were disregarded, as they are related to metals: "830220 Castors: with mountings, of base metal" and "970300 Sculptures and statuary: original, in any material". Both would have been matched to marble. The following metals were not exported from Myanmar to China in 2011 and are therefore matched by the difference in export value in the first year after 2011 in which there are exports: Chromium (matched in 2012), Baryte (2013), Limestone (2014), and Gypsum (2017).