Essays in Spatial Economics

by

Antoine Boris Levy

B.S., Universite Paris-I Pantheon-Sorbonne (2012)

M.S., HEC Paris (2016)

M.S., Paris School of Economics (2016)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2022

Author . . . .

Department of Economics

May 13, 2022

Certified by . . . .

Arnaud Costinot

Professor of Economics

Thesis Supervisor

Certified by . . . .

Ivan Werning

Professor of Economics

Thesis Supervisor

Accepted by . . . .

Abhijit Banerjee

Ford International Professor of Economics, Department Committee on Graduate Theses

Essays in Spatial Economics

by

Antoine Boris Levy

Submitted to the Department of Economics
on May 13, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis is composed of three essays, showing how nationwide economic causes exert distinct local and aggregate effects across regions, depending on the geographic distribution of exposure to these common shocks, and on spatial interactions between locations.

The first chapter, building upon administrative data covering the universe of dwellings in France, documents the presence of home bias in investment (a negative effect of distance for individual investors' lumpy portfolio allocation decisions). I explore its consequences for the equilibrium supply of housing, in a spatial equilibrium framework combined with a frictional portfolio choice. Using quasi-experimental evidence from a location-specific French investment tax credit targeted at individual landlords, I evidence a substantial causal impact on transactions, new construction, investor returns, and inwards migration. Long-distance individual investor involvement rises in treated cities, and the policy has stronger effects in locations more open to outside capital.

The second chapter, in collaboration with Jacob Moscona, studies how exogenous differences in local population density lead regions to specialize in different kind of manufacturing industries. We show theoretically and empirically that a country's economic geography – in particular, the distribution of population across space – is an important source of comparative advantage, as countries with higher population-weighted population density specialize in sectors that benefit from agglomeration. After estimating substantial variation within the US in the extent to which manufacturing sectors sort into dense locations, we find that countries with higher population-weighted density disproportionately export in sectors with high "density affinity".

The third chapter explores electoral behavioral with regionally differentiated exposure to common campaign pledges. Using quasi-random spatial variation across municipalities, and an instrumental variables strategy exploiting formulaic real estate assessments established in the 1970s, I show that a promise to repeal a broad-based housing tax accounted for a substantial share of Emmanuel Macron's electoral success in the 2017 French presidential election. In high-frequency data, the timing of the promise coincided with a significant increase in voter information search, in Macron's polling intentions, and in his market-based predicted chances of victory. The results evidence the crucial role of spatial distributive policies, even in elections marked by ideological polarization around non-economic issues.

Thesis Supervisor: Arnaud Costinot
Title: Professor of Economics

Thesis Supervisor: Ivan Werning
Title: Professor of Economics

Acknowledgments

I owe more thanks than I can describe to Arnaud Costinot and Ivan Werning. I came to MIT thinking of myself as a budding international macroeconomist, and of Ivan and Arnaud as the finest economists in this field to take advice from. My interests gradually shifted, from theory to applied research, and from the borders of countries to the less formal – but often no less consequential – boundaries that divide regions within them. Nevertheless, because they are economists in the broadest possible sense, Arnaud and Ivan, each in his own unique style, remained the most helpful advisors and acute critics of my work I could ask for. For their unfailing support, even when things sometimes seemed to fall apart, I will remain grateful. Jim Poterba's sharp eye for novel questions in public finance, and his indulgence as I wandered from one topic to another, were also of tremendous importance for me.

I was fortunate to count on the support of numerous mentors along the way. I would not have become an economist had I not witnessed firsthand the brilliance and creativity of Daniel Cohen. His kindness was also the reason why, amid unexpected pandemic disruptions, I found a temporary safe haven at the Paris School of Economics, boulevard Jourdan, for more than a year. I am deeply indebted to the intellectual rigor of many of those who have provided me with informal advice over the years: Daron Acemoglu, Treb Allen, David Atkin, Abhijit Banerjee, Agnes Benassy-Quere, Olivier Blanchard, Sylvain Catherine, Dave Donaldson, Jason Furman, Peter Ganong, Pierre-Olivier Gourinchas, Jon Gruber, Xavier Jaravel, Thomas Piketty, Alp Simsek, David Thesmar, and many others. I am also thankful to countless institutions, universities, and people, whose seminar comments, financial support, or data provision were critical to this thesis; notably the Paris School of Economics, Dartmouth College, the Lincoln Institute for Land Policy, the Mercatus Center, the CEREMA, the MIT Shultz fund, and the Sloan coffee lounge. The administrative staff at MIT Economics (in particular Julia Martyn-Shah, Shannon May, Eva Konomi, Gary King, Jonathan Brownrigg, Ruth Levitsky, and Kim Scantlebury) deserve special thanks for their tireless efforts, especially during the academic job market.

The many people I met in the Cambridge bubble filled these PhD years with excitement: Jonathan Amar, Maxime Amram, Zachary Bleemer, Natasha Carnel, Oguzhan Celebi, Antoine Dedieu, Mayara Felix, Joel Flynn, Paul Fontanier, Isadora Frankenthal, Joe Hazell, Clemence Idoux, Andrea Manera, Jeremy Majerovitz, Jacob Moscona, Jean Pauphilet, Charles Serfaty, Jad Sassine, Karthik Sastry, Evan Soltas, Jeremy Toledano, Joonas Tuhkuri, Pierre-Luc Vautrey, Olivier Wang, Sean Wang, Jonathan Zonabend Baravir. Some of them prospective or current co-authors, dinner and cocktail companions, squash opponents, cohort mates, office mates, flatmates; all of them friends. This thesis owes a lot, directly and indirectly, to conversations with them - it is hard to say, in particular, if I learnt more from my co-authorship with Jacob or from our friendship. Elsewhere in the world, my closest friends have always been there to cheer me up and to remind me of the correct priorities in life: Baptiste, Quentin, Simon, Benjamin, Thomas M, Thomas L, Paul-Adrien, Hannah, Camille, Florence, Charles, Karen, Esther, Hamouda, Sixtine, Mathieu, David, and many others.

Completing this PhD would not have been possible without the constant and loving support of my family - above all, my father Marc and my sister Judith, who stood by me in good and bad times, as they always have; but also my grand-parents, my aunts and uncles, my cousins, Thibaud, Norah, Corinne, Rene, Thomas and Victor. I have a special thought for Yannick - the first in my family to have gone through a PhD, whose advice one evening in Berkeley, in 2016, convinced me to follow his path.

With economists, co-authors, and cohort mates, one talks about research. With friends, family and loved ones, one contemplates the rest. Only one person, however, is all of this at once, and could understand it all: the linear regressions and the non-linear progression; the home bias in my data and the homes we created for us in Toulouse, Paris, or Cambridge; the simplistic mobility decisions of our models, and the complex ones made in our own life. Your brilliance, your rigor, your talents as a researcher continue to amaze me since the first time we met. Your humor and your joy illuminate every day I am lucky enough to spend with you. I hope to forever deserve the love and trust that you placed in me – this year more than ever. Thank you, Mathilde.

This dissertation is dedicated to my mother. While I often wonder if she would have expected me to follow the academic path, I hope she would be proud that I did so to the fullest. I miss her every day.

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

Chapter 1

Housing Policy with Home-Biased Landlords [1]

The landlords, like all other men, love to reap where they never sowed.

A. Smith, An Inquiry into the Nature and Causes of the Wealth of Nations, I, 6.

1.1 Introduction

A limited supply of housing constrains the growth of cities (Glaeser and Gyourko, 2018). In particular, the scarcity of rental accommodation – a key driver of spatial mobility – severely limits migration towards desirable or productive areas. In response, governments in advanced economies have devised a wide array of place-based interventions to alleviate housing burdens for renters in expensive real estate markets. These regulatory and fiscal tools must balance two conflicting goals: making rental housing affordable for tenants, while letting its provision remain profitable to landlords. Navigating this trade-off, in turn, requires understanding the constraints that govern investment behavior across local housing mar-

[1] I thank Gabrielle Fack, Etienne Lehmann, Yuhei Miyauchi, John Sturm, Bill Wheaton, and seminar participants at the MIT Trade tea and Macro and Public Finance lunches, the OFCE, the Paris School of Economics, the Toulouse School of Economics, the Wharton School, Stanford GSB, Berkeley Haas, the San Francisco Federal Reserve, the FDZ Ruhr, the London School of Economics, CREI, for helpful discussions and feedback. I am indebted to countless officials at the French Treasury, the Housing Secretariat, and the Defence Ministry, for their help in accessing the underlying data. This work also benefited from the generous support of the George and Obie Shultz fund at MIT, and public grant ANR-10-EQPX-17 of the French National Research Agency (ANR) as part of the "*Investissements d'avenir*" program.

kets. Do financial and distance barriers affect whether individual savers decide to become landlords, and where? To what extent do tax incentives shape the geography of new housing investment? And how do spatial frictions determine the way these policies impact the portfolio allocation of investors, the supply of housing, and the mobility of renters?

Most studies of the forces that lead economic activity to cluster in space document channels operating in the markets for labor (Marshall, 1890), knowledge (Jacobs, 1969), or final goods (Krugman, 1991). Paradoxically, on the other hand, public policies shaping the economics of cities concentrate on housing investment, where academic research generally abstracts from space and assumes a perfectly elastic supply of "absentee" capital. This paper provides a possible reconciliation between the policy focus on housing production, and the theoretical argument for endogenous agglomeration. It illustrates empirically and conceptually that landlord ownership decisions are biased towards rental properties close to their own residence. Notwithstanding Adam Smith's view, landlords love to reap precisely *where* they sowed. The home bias of these immobile individual savers, who supply imperfectly elastic capital across space, affects the provision of rental housing. It makes it relatively cheaper for tenants to agglomerate in cities with high proximity to affluent savers – a "market access" term, in the terminology of Redding and Rossi-Hansberg (2017), that operates through frictional capital markets. Critically, spatial frictions also shape the response of housing markets to tax incentives. While individual residential investment, like business capital formation (Zwick and Mahon, 2017), responds to tax policies, its specific spatial pattern has implications for the mechanisms and heterogeneity at play. Using quasi-experimental evidence from a French investment tax credit, I demonstrate how, consistent with a stylized model, the geography of home-biased landlords influences the heterogeneous effects of place-based rental subsidies across eligible locations, and regulates their impact on the mobility of capital and people in the economy.

To investigate constraints to the supply of rental services by individual savers, and analyze the dynamics of production, sales, and occupancy for the rental housing stock, I assemble register data on the universe of close to thirty million primary residences in France. I also avail myself of exhaustive administrative records for all real estate transactions, building permits requests, and individual landlord tax returns, along with detailed data on land use, social housing, mobility, and employment, covering France's 35,000 municipalities over the 2010-2020 decade. The paper then proceeds in three main steps.

I begin by documenting three facts illustrating that rental housing investment is a highly frictional venture. Reminiscent of common paradoxes in international finance (Feldstein and Horioka, 1980; Lucas, 1990; French and Poterba, 1991), rental capital flows within a country exhibit significant spatial segmentation, evidenced by three regularities: portfolio lumpiness, return dispersion, and local bias in ownership.

First, wealth survey microdata illustrate that rental property is mostly held in lumpy quantities by undiversified and mostly immobile landlords. Far from the ideal view of a common diversified housing market portfolio, each of them tends to operate a small number of units, a limited scale suggestive of physical indivisibility and financial constraints. Next, I show there exists substantial dispersion in rent-to-price ratios, both across cities and among individual units within narrow locations. That variation in capitalization rates differs from their hypothetical equalization to a unique user cost, as expected in integrated capital markets. While it could stem from differences in risk or expected capital gains, the additional fact that rent-to-price ratios decline with the distance of a location from potential investors suggests that heterogeneous access to investors matters in the market for residential capital. Lastly, an analysis of more than seven million rental properties in France demonstrates that investment linkages between cities (the probability of ownership of rental assets in one of them by investors from the other) decay rapidly with their physical distance and estimated social connections. The geography of rental asset ownership displays a substantial "housing home bias" of landlords for dwellings near their own residence location.

Together, these regularities make investment frictions essential in shaping the allocation of capital across markets, and, in turn, of mobile workers across space. Building upon and rationalizing these facts, my second contribution is a stylized model where barriers to cross-city capital mobility constrain new housing production, and influence its response to tax policies. While the imperfect mobility of *workers* (Kline and Moretti, 2014a) and *firms* (Fajgelbaum et al., 2019) has been the topic of a long spatial literature, distance barriers to the allocation of *housing capital* by individual investors remain unexplored.

In the model, renters, as in a standard Rosen (1979)-Roback (1982) framework, move across locations to arbitrage real wage differentials. Housing producers build new units on local land, but do not lease them directly; rather, they sell them outright to investors, who are the sole providers of rental services. I depart, however, from the standard assumption that deep-pocketed, well-diversified, and "absentee" landlords require a uniform rate of return on their rental properties, a perspective under which the free flow of capital would equalize rent-price ratios across locations. Instead, these dwellings are owned by overlapping generations of spatially static individual savers, whose location-specific rental asset demand reflects relative returns, operating costs, and idiosyncratic frictions or preferences for operating specific residential projects.

Because these subjective investment frictions increase with distance, landlords prefer to operate rental dwellings close to their own residence, willingly accepting lower gross returns on nearby properties. Their lumpy investment decisions and responsiveness to return differentials shape housing affordability in destination markets. As in the data, the ownership of the rental stock exhibits a model-implied gravity structure, and areas located close to an ample supply of investible capital have lower rent-to-price ratios, due to their higher "investor portfolio access". Alongside long-studied local externalities (Duranton and Puga, 2004), the capital-intensive but frictional production of rental services thus creates a novel agglomeration force for the location choice of mobile workers – who are overwhelmingly tenants. This frictional investment behavior also pins down spatial heterogeneity in the tax responsiveness of rental housing supply: locations more open to outside capital draw from a broader and more distant pool of individual investors when landlord subsidies are implemented. With potential landlords reluctant to invest far from home, the market for rental assets is segmented, and the upwards-sloping supply of capital less than perfectly elastic across space: the incidence of place-based housing subsidies partly benefits infra-marginal savers in the form of higher returns.

In the third part of the paper, I present direct quasi-experimental evidence that, consistent with the model, individual landlords not only respond sharply to housing tax policies, but also provide an imperfectly elastic capital supply across space. To quantify the causal response of landlord behavior to place-based incentives, I study a French tax credit targeted at the provision of new rental housing by individual investors. The scheme, known as the *Pinel* law, offers taxpayers a substantial personal income tax reduction (of up to 21 percent of the purchase price) if they buy and lease newly built properties in targeted high-cost areas, conditional on respecting affordable rent levels and tenant income ceilings. I take advantage of variation across comparable municipalities in the place-based eligibility of buy-to-let subsidies after their introduction in 2014, and of their partial removal in 2018, when the *Pinel* incentive was discontinued in some cities as a byproduct of nationwide budget cuts. Leveraging these reforms, I combine difference-in-differences around entry and exit for similar treated and control locations, to test model predictions on the consequences of place-based subsidies.

First, tax policy causally affects the allocation of individual *investment* across space. Developer sales of eligible newly built dwellings to individual investors increase by 18-20 percent each quarter in targeted locations for the full duration of the subsidy (4 years), with the effect operating at both the intensive and extensive margins. The implied user cost elasticity of investment ranges from 1.3 to 2, depending on an estimated net-of-rent-discounts present value of the incentive between 9 and 15 percent. Far from merely accelerating sales of existing developer inventory, the policy triggers entirely new construction projects, increasing local housing starts and land demand. The composition of dwellings also shifts, as the share of multi-family buildings and social housing increases, and urban sprawl rises. Going beyond the difference-in-differences design, in the longer-run, nationwide rental property purchases closely co-move with the time-varying generosity of government subsidies, as predicted by the model. Total new home sales also display stark time-bunching in anticipation of eligibility deadlines for buy-to-let tax credits, and abnormal end-of-fiscal-year spikes, hinting at the salience of taxation for aggregate housing investment.

Next, I provide direct evidence that the effects are consistent with the spatial bias mechanism of the conceptual framework. First, in contrast to common findings for real estate subsidies, but in line with an imperfectly elastic capital supply across space, there is only incomplete capitalization of the incentive in new unit prices. The estimated price effects only amount to between one third and one half of the net-of-rent-discount present value of the landlord subsidy. The policy thus raises post-tax returns to rental property investment in treated locations to attract additional capital. Next, as predicted by the model's extensive margin of portfolio reallocation, the scheme operates by drawing in more distant marginal investors to treated locations, somewhat offsetting home bias in rental investment. Finally, I show evidence that the scheme displays stronger effects in locations more open to outside ownership, confirming that the degree of *capital* mobility has far-reaching consequences for the provision of housing across places. The policy results in a 2 percent long-lived expansion in the local housing stock, with a larger effect in locations characterized by a more spatially dispersed ownership.

Last, the portfolio choice of investors entails a downstream reshuffling of *people*: spatial capital reallocation induces labor mobility. Contrary to the existing literature on housing policy, I find no evidence of crowding-out of non-subsidized housing. As a consequence of this limited crowding-out, in the medium-run, the scheme triggers a positive response of geographic mobility towards towns eligible to the individual landlord subsidy. It leads to increases in local population and employment, mediated by a sharp upwards jump in inwards migration (notably of middle- and lower-income tenants) towards treated cities.

Related literature This paper evidences the relationship between the spatial allocation of economic activity, the provision of housing, and the geography of wealth. It builds on three main strands of investigation: on the implications of housing supply for the spatial economy; on the heterogeneous consequences of tax incentives for investment behavior; and on fiscal policies in housing markets.

First, I present a new channel for the role of housing in the spatial allocation of economic activity, operating through capital market frictions. Ganong and Shoag (2017) and Glaeser and Gyourko (2018) show that zoning regulations limiting local housing supply elasticities constrain mobility and growth, while Monte, Redding, and Rossi-Hansberg (2018) demonstrate that openness to commuting can circumvent a limited local land supply. In this literature, however, the decisions of "absentee" landlords generally take a back seat under the assumption of perfect capital mobility. I provide evidence on the specific role of private rental provision for triggering labor mobility, and explore how capital market fluidity across space matters for rental affordability. Recent work on rental housing implies that landlord heterogeneity (Greenwald and Guren, 2019) and access to finance (Gete and Reher, 2018;

Reher, 2021) may generate aggregate supply curves that are upwards-sloping in returns. In my paper, systematic geographic variation in landlord user costs makes the supply of rental capital depend on physical proximity to investors, potentially misallocating the provision of dwellings.[2] This local bias also has distributional consequences. Due to the spatial frictions in residential investment, location-specific subsidies involve trade-offs between redistribution and incentives across places.[3] In the same way that limited *labor* mobility makes place-targeted measures fall upon infra-marginal *workers* (Kline and Moretti, 2014b; Gaubert, Kline, and Yagan, 2021), here, frictions to the allocation of *capital* across space entail a partial incidence of place-based policies on nearby *savers*.

Second, I study how frictions alter the sensitivity of investment to tax policy, a public and corporate finance question dating back to Hall and Jorgenson (1967). Zwick and Mahon (2017) and Ohrn (2018) exploit industry variation in depreciation allowances to identify the tax responsiveness of firm capital formation. Using quasi-experimental variation in *legal eligibility* (rather than differential exposure) across *locations* (rather than sectors), this paper estimates the response of *residential* investment (rather than business equipment) by *individual savers* (rather than firms).[4] Criscuolo et al. (2019) and Suárez Serrato and Zidar (2016) also analyze spatially varying tax incentives, but focus on firm capital and labor demand.[5] In this literature, heterogeneous corporate responses to taxation provide indirect evidence of firm-level financial constraints. Here, by treating landlord investment as a geographically differentiated choice, I demonstrate that heterogeneous responses to tax credits across places inform us on spatial frictions in individual portfolio allocation decisions.

Finally, I contribute to a growing body of work exploring the causal impact of government involvement in real estate markets, heretofore mostly dedicated to housing *consump-*

[2]The spatial segmentation of housing ownership that I document meets a sizable body of evidence on international (Portes, Rey, and Oh, 2001) and domestic (Coval and Moskowitz, 1999; Hong, Kubik, and Stein, 2008; Chen et al., 2010) home bias in asset holdings. The consequences of local bias for rental affordability also relate to the hypothesis that proximity to capital (Kuchler et al., forthcoming), or limited local competition in asset markets (Hong, Kubik, and Stein, 2008), have downstream effects on returns.

[3]Other studies, e.g. Favilukis, Mabille, and Van Nieuwerburgh (2019), show the potential for redistribution through housing policies.

[4]Sommer and Sullivan (2018) study the response of residential investment to nationwide tax policy changes in a full-scale calibrated model with no spatial dimension.

[5]A recent body of work studies the Opportunity Zones created by the Tax Cuts and Jobs Act of 2017 in the United States (Freedman, Khanna, and Neumark, 2021; Chen, Glaeser, and Wessel, 2019; Kennedy and Wheeler, 2021), mostly finding null or modest effects of this large-scale place-based policy on local outcomes, with the exception of Arefeva et al. (2021).

tion – rather than *investment*.[6] The buy-to-rent subsidies I study provide tax rebates to investors in rental dwellings, often conditional on affordability criteria. In spite of the growing use – and budget costs – of subsidized privately-owned rental property, its causal effect on investment behavior, spatial asset ownership, and tenant mobility remains largely unknown, along with its distributional consequences. In the United States, the Low-Income Housing Tax Credit (LIHTC), whereby the federal government shoulders part of the cost of affordable housing construction by corporate developers, has attracted attention as its fiscal burden rose in recent years (Baum-Snow and Marion, 2009; Diamond and McQuade, 2019). The French *Pinel* law shares similarities with the LIHTC: both provide tax benefits to newly built dwellings in exchange for rent moderation. However, relative to the LIHTC, the *Pinel*'s open-ended nature, the nationally pre-determined eligibility of locations, and its targeting of individual taxpayers, further our understanding of residential investment responses, capital market segmentation in housing production, and agglomeration effects across income groups.[7]

[6] Since Poterba (1984)'s seminal contribution, other papers have explored tax subsidies to home-ownership, from mortgage interest deductions (Hilber and Turner, 2014; Gruber, Jensen, and Kleven, 2021) to home-buyer tax credits (Berger, Turner, and Zwick, 2020; Carozzi, Hilber, and Yu, 2019) or stamp duty exemptions (Best and Kleven, 2018). Studies of the rental market have been mostly confined to tenant-side support for housing consumers, evaluating vouchers (Eriksen and Ross, 2015; Collinson and Ganong, 2018) or cash allowances (Brewer et al., 2019; Fack, 2006). Olsen and Zabel (2015) reviews the literature on housing market policies; see Metcalf (2018) for a survey focused on housing affordability.

[7] Early cross-sectional studies of the LIHTC found close to full crowding-out effects on un-subsidized housing (Malpezzi and Vandell, 2002; Sinai and Waldfogel, 2005; Eriksen and Rosenthal, 2010). Because the *Pinel* law is targeted at individual – rather than corporate – investors, it has direct implications for the respective role of returns, tax incentives, and spatial frictions in individual landlord decisions. Furthermore, its effects on the local and national housing markets are likely to be distinct from the LIHTC: sales claiming the *Pinel* rebate represented close to half of overall sales of new multi-family units in France in 2016-2018; broad swaths of the population are eligible as tenants; and the mechanism is "open-ended", while the LIHTC relies on fixed allotments and administrative approvals. Empirically, novel administrative data allow me to examine the role of home-bias at the level of individual units, landlords, and local communities. Finally, while studies of the LIHTC generally rely on *ex post* variation in distance to, and timing of, actual projects implemented, my strategy instead exploits *pre-determined* spatial targeting and sharp eligibility changes over time, without assuming "ring"-based treatment intensity decay. This paper also complements evidence on the inclusionary zoning participation of developers (Soltas, 2020), but focuses on the intensive margin of housing production, and the allocation of their residential asset portfolio by individual savers.

1.2 Home-biased landlords: three motivating facts

As of 2019, forty percent of households in France rented their homes, with close to two thirds of these tenants residing in about seven million privately provided rental units.[8] Leveraging register data on all dwelling units in France, exhaustive transaction deeds for the entire country, as well as individual asset holdings and income tax returns microdata, this section documents the financial and geographic frictions constraining residential investment by individual landlords. First, most private rental dwellings are leased at low scale by affluent but undiversified "mom-and-pop" investors. Second, properties in the data exhibit highly dispersed rental yields, both across towns and within the same location, with lower rent-to-price ratios in locations located closer to potential investors. Third, landlords display a local bias for nearby properties, with rental ownership probability fast decaying in the distance between investors and their assets. These findings motivate a framework where the location of savers, and their investment behavior sensitivity to returns and spatial frictions, take center stage in the allocation of dwellings and people across space.

Fact 1: Residential investment is lumpy and operated at low scale by undiversified households I first use household wealth survey data to document the lumpy, concentrated, and undiversified nature of rental housing ownership.[9] Linking asset ownership information to landlord status evidences several relevant features of the French rental property market. First, as evidenced in panel (a) of figure 1-1, most landlords in France are so-called "mom-and-pop" investors, who own two or fewer housing units. Rental property is highly lumpy and frictional: two thirds of investors own only one unit, and, as panel (b) of figure 1-1 illustrates, close to seventy-five percent of all dwellings are owned by landlords possessing two units or fewer.[10] Second, the close to 3 million landlord households are substantially wealth-

[8] Aggregate figures are computed from the *Base Logements*, a combination of the Housing survey and Census conducted by France's national statistical agency (*INSEE*). The remaining third of tenants live in publicly provided "social" housing, offered by the national governments at deep discounts against market rents. In the United States, c. 36 percent of households were renters as of 2019Q4, according to Census data.

[9] Appendix 1.C provides additional information on the restricted-access microdata from the 2018 household wealth survey (*Enquete Patrimoine*) used in this subsection.

[10] The implied total number of units leased by households in the wealth survey (4.3 million units) is not equal to the total number of rental dwellings owned by private investors in the individual dwellings data (around 7 million in metropolitan France as of 2018). The discrepancy is likely explained by three main factors: non-response, lack of representativeness, or measurement error in the survey data (e.g. when investors

ier than both renters and non-landlord owner-occupiers,[11] as shown in panel (a) of figure 1.E.2. The median wealth of landlords is above the 99$^{\text{th}}$ percentile of net worth for renters: this entails a clear wealth segmentation between the population of moderate-income renters and the more well-off group of owners of the housing stock.[12] Finally, panel (b) of figure 1.E.2 computes the portfolio share of wealth in rental dwellings among landlords. The median share is 52 percent, and rental assets exceed 30 percent of wealth for three quarters of owners. For most landlords, the value of rental property constitutes a substantial proportion of their assets, indicating indivisibilities and frictions to risk-sharing and portfolio diversification. Overall, similar to the way financial frictions can lead to a misallocation of capital in the production of tradable goods (Aghion and Bolton, 1997; Moll, 2014), operational barriers prevent rental housing from being operated at large scale by the most efficient landlords. The lumpy, undiversified, and individual nature of rental ownership makes the spatial distribution of ownership relevant to aggregate rental productivity and housing affordability.

state ownership of one dwelling but actually mean a multi-family building with several units); ownership by institutional investors, which is limited and represents between one and two percent of the overall private residential rental housing stock in France, according to INSEE; and foreign ownership of rental housing. However, since the wealth survey data are mostly used to document investor characteristics, rather than properties of the overall housing market, this partial lack of coverage is not consequential for the regularities evidenced here.

[11] An overwhelming majority of landlords in the data (close to 90 percent) are themselves owner-occupiers.

[12] The median net worth of landlords is about €500,000 in 2018, while the corresponding figure for non-landlords owner-occupiers is €217,000, and €10,000 only for renters. Appendix 1.E.2 provides some additional demographic information on the characteristics of owners of rental dwellings, who are older and have higher income than both non-landlords owner-occupiers and renters.

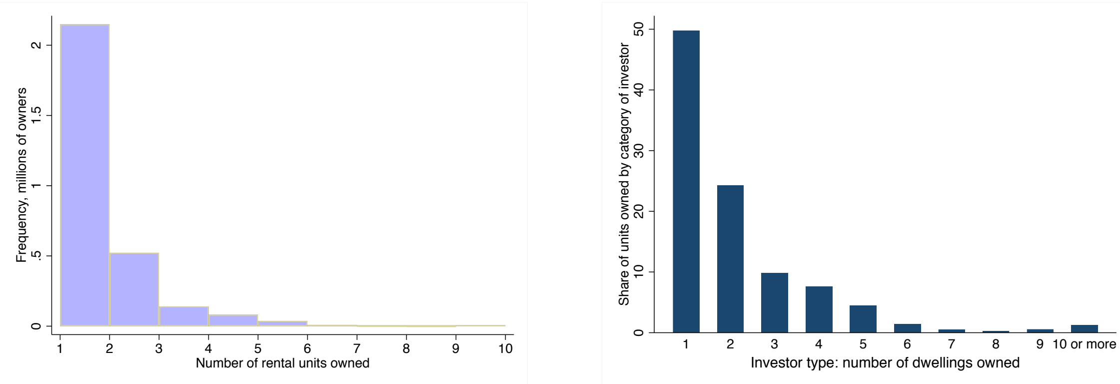

(a) Share of landlords by size category
(b) Share of units by landlord size category

Figure 1-1: **Landlords and the lumpy ownership of rental housing**

Panel a plots the distribution of landlords according to the number of rental units they state to own, using matched household-assets data from the wealth survey. Panel b plots the distribution of units owned by landlord size category. Data are computed at the household level, from the *Enquete Patrimoine* 2017-2018.

Fact 2: Rates of return are unequal across space and units, and decrease with proximity to investors I next quantify the large dispersion in contemporary capitalization rates (the ratio of current rents to housing prices) across local housing markets, as well as within them, to provide further evidence of frictions in the allocation of residential investment. I first use town-level 2018 data on rents[13], and calculate the median acquisition price per square meter for multi-family units in exhaustive transactions deeds data for the three years immediately surrounding 2018, to compute rent-price yields at the municipality level.[14] Although rents and sales prices are highly positively correlated across cities, rent-price ratios still vary widely across space, as shown in panel (a) of figure 1-2. Capitalization rates are not scattered at random across locations; instead, as shown in panel (a) of figure 1.G.2, I find systematically lower rent-price ratios in cities with higher prices.

Gross yield dispersion across municipalities, however, may be driven by variation in the

[13]No systematic panel data on rents across locations exists in France. Only one official cross-sectional measure of representative 2018 rents per square meter in multi-family apartments across municipalities, is made available by the French Housing secretariat.

[14]Municipalities, or *communes*, are the smallest French administrative unit, often comparable to US ZIP codes. Their average area is 15 square kilometers (average population of 1,850), although their size varies widely: the largest one, Paris, counts 2.1 million residents. The transactions deeds data are described further in appendix 1.C. To limit measurement error concerns, I focus here on municipalities with more than 11 transactions for multi-family units over the three years period. Appendix figure 1.G.3 verifies that the results are not driven by the selection of larger municipalities, providing similar evidence for all towns with no restriction on the number of transactions.

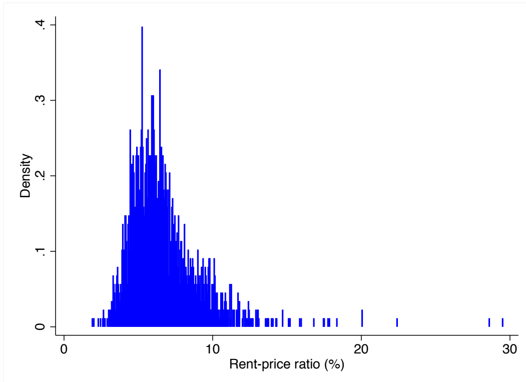

(a) Rent-to-price: across municipalities

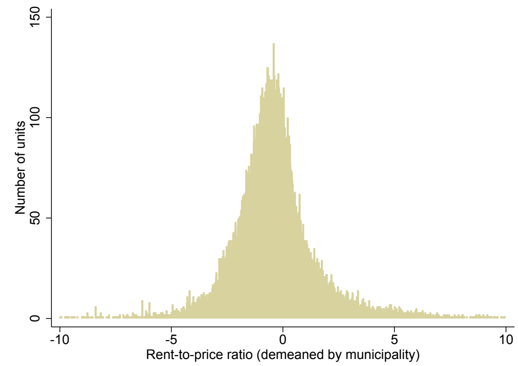

(b) Rent-to-price, buy-to-let units: within municipalities

Figure 1-2: **Dispersion in rent-to-price ratios**

Figure 1-2 documents the dispersion in rent-to-price ratios across and within municipalities. Panel a plots the distribution of rent-to-price ratios across municipalities, using average rents per square meter for multi-family units in 2018, and median acquisition prices for multi-family units from 2017 to 2019 in transaction deeds data (*DV3F* database). Panel b plots the distribution of re-centered (by municipality-year level mean) rent-to-price ratios across rental units purchased under the *Pinel* subsidy scheme, for municipalities with more than one unit, using individual tax returns (*POTE* database) for all purchases made in 2016 under the *Pinel* scheme in 2016 (see main text).

average vintage of current local leases, or differential expectations of capital gains and rent growth. Moreover, rentals may differ from recently transacted units due to selection bias (Eichholtz et al., 2021), making cross-city rent-price-ratio comparisons difficult.[15] To illustrate the large residual variation in rental returns even after controlling for expected local dynamics, I turn to a measure of the dispersion of within-location-year unit-level gross rental yields, using confidential individual tax returns. To do so, I exploit the fact that households who purchase a buy-to-let unit under a subsidized housing scheme in France must fill out a "buyer's commitment form" stating the purchase price and initial rent of the unit.[16] I

[15]In a standard Campbell-Shiller decomposition (Campbell and Shiller, 1988) decomposition, current rent-price ratios reflect expected changes in future rents and capital gains stemming from changes in discount rates (Campbell et al., 2009). Within-lease rent changes are heavily regulated in France. For the same tenant, the maximum annual increase is limited to the yearly evolution of a nationwide official index (the "*Indice de Reference des Loyers*" or *IRL*) computed on the basis of the overall CPI (excl. rents and tobacco). In case of a tenant change, the rent can increase by more than the *IRL*, except since 2012 in so-called "tense" property markets covering large cities, where new-lease rent increases are also limited by the *IRL*, unless the property has been vacant for more than eighteen months or has undergone substantial improvement works. As regards capital gains, a standard model with a representative owner and mean-reverting prices suggests that higher future capital gains should be correlated with *lower* current prices, leading to counter-factually *higher* rent-price ratios in *low* price locations (Sivitanides et al., 2001; Liu, Wang, and Zha, 2019).

[16]Buyers must fill this form, known as *schedule 2044EB*, when they first file their individual income tax returns in the year of the purchase, in order to benefit from the conditional rental investment tax credit. The

compute rent-to-price ratios for close to 65,000 housing units purchased under the *Pinel* subsidy scheme in 2016. Demeaning rent-price ratios by a municipality-year-level average, the distribution of re-centered rates of return in panel (b) of figure 1-2 still exhibits substantial dispersion in returns within a location.[17] While maximum rents per square meter are regulated by a ceiling under the *Pinel* subsidy scheme (see *infra*, section 1.4), both actual rents per square meter and purchase prices vary widely within towns. Similar to the cross-city evidence, panel (b) of figure 1.G.2 shows rent-to-price ratios are lower in higher-price units, even after accounting for municipality-level fixed effects.[18]

Moreover, beyond the large degree of dispersion, I show that the rent-price ratio in a location is negatively correlated with the proximity of the city to potential investors. Using income tax data aggregated across France's close to 300 commuting zones, I first compute a proxy for local wealth W_i in location i as the sum of local capital income earned by residents of i, $W_i \propto \tilde{W}_i = \sum_{z \in i} y_z^K$.[19] I then calculate an empirical proxy for the "investor portfolio access" of location j, an index for its physical proximity to investors, as the sum of wealth in all locations i, inverse-weighted by a function of their bilateral distance to j, d_{ij}: $IPA_j = \sum_j e^{-d_{ij}} W_i$.[20] Figure 1-3 evidences that a commuting zone's rent-to-price ratio systematically and negatively correlates with investor proximity: cities that are physically closer to savers exhibit lower equilibrium rent-price ratios.

form includes the purchase price of the unit, its location, the purchase date, as well as the initial rent at which it will be leased to the tenant.

[17]City fixed-effects only account for 13 percent of the overall variance in rent-price ratios in the data consisting of buyers' commitment forms in individual tax returns.

[18]Asset-level variation in gross rental yields even within narrow locations accords with the idiosyncratic risk in Giacoletti (2021).

[19]Local wealth data are not directly available, because the wealth survey does not record precise household location information. I use aggregate taxable capital income from administrative data at the commuting zone level, as a proxy for the ranking of locations by wealth. In the spirit of the wealth inequality literature's common use of the "income capitalization" method (King, 1927; Stewart, 1939), under the extreme first-order approximation of a single capitalization factor, log wealth in a location would be directly proportional to log capital income in a location (Chodorow-Reich, Nenov, and Simsek, 2021).

[20]The choice of aggregation at the CZ level relates to the computational feasibility of measuring all possible distances across location pairs: for close to 300 mainland CZs, the number of bilateral pairwise distances to estimate is about 45,000, but it would rise exponentially when estimating the entire network of pairwise distances across France's 35,000 municipalities. Anticipating on the predictions of the framework in section 1.3, I take an inverse exponential function as my baseline approach to inverse-distance weighting of access to wealth.

Figure 1-3: **Rent-price ratios and investor proximity**

Figure 1-3 plots a binned scatter plot of the relationship between the rent-to-price ratio, $\frac{R_j}{P_j}$, and (the log of) the estimated "investor portfolio access" $IPA_j = \sum_i W_i e^{-d_{ij}}$ on the x-axis, where d_{ij} is the bilateral haversine distance between CZs i and j, and W_i is a measure of total capital income in i. The figure is estimated across France's 293 mainland commuting zones with available data and town-level rent-price ratios within a CZ are aggregated after weighting by the number of transactions over 2017-2019.

Taken at face value to embody variation in real returns to housing investment,[21] the dispersion in rent-price ratios itself (figure 1-2) suggests a "Lucas paradox" of residential investment: why does it not flow from low to high rate-of-return locations and units? While merely suggestive of the presence of distance frictions, the decreasing relationship of rent-price ratios with proximity to investors (figure 1-3) offers one potential explanation, further explored in my stylized model: the spatial segmentation of rental investment may entail an inelastic supply of imperfectly mobile capital, which fails to equalize returns across space.

Fact 3: There is a "home bias in homes" and landlord investment decays with distance Lastly, I document spatial constraints to rental property investment directly from exhaustive register data on the universe of rental properties in France and their owners. In the absence of any distance barriers in the rental market, the residence of landlords owning property in any given city would follow the same spatial pattern as the overall landlord population, as

[21] As in the international finance literature on investment frictions (Caselli and Feyrer, 2007), observed returns may be subject to measurement error, which can shroud the true variation across locations and units. I take the substantial observed dispersion in rental rates, and their systematic decreasing pattern with investor portfolio access, to be a stylized regularity consistent with impediments to the flow of capital across cities in the rental market.

a byproduct of portfolio diversification. In practice, however, this null hypothesis of perfect diversification is strongly rejected. Housing investment exhibits substantial home bias towards areas close to the residence location of the landlord-investor. The count of units $\mathcal{C}_{ij}$ in rental market j owned by investors living in i decays substantially with bilateral geographic and social distance d_{ij} between the two cities.[22]

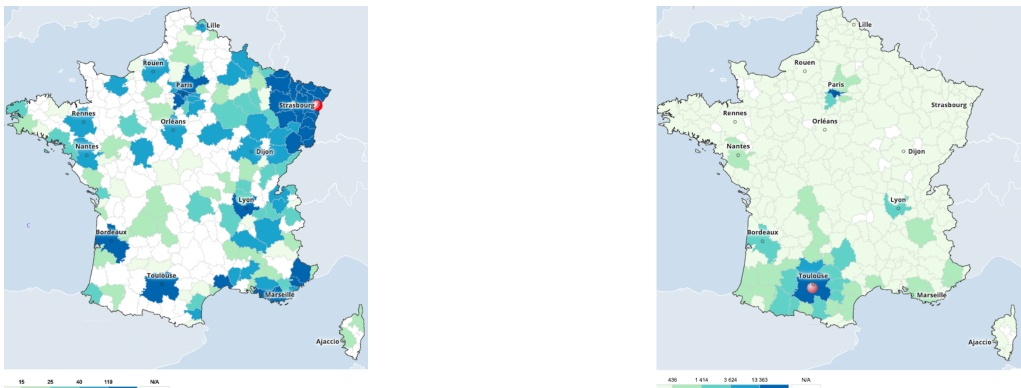

(a) Ownership of Strasbourg rental units (b) Ownership of Toulouse rental units

Figure 1-4: **The spatial pattern of rental asset ownership**

Figure 1-4, panel a (resp. b), plots, for each of 304 commuting zones in mainland France and Corsica, the number of units owned in CZ j : Strasbourg (resp. j : Toulouse) by landlords residing in any CZ i (C_{ij}), including in the same commuting zone. CZs are ordered into 5 classes of ownership intensity of destination CZ units by origin CZ investors, shaded from lightest to darkest, according to the Jenks-Fisher natural breaks optimization routine. In accordance with data disclosure agreements with the French Treasury, only CZs where landlords own more than 11 units in the destination rental market are shown. Statistics are computed from individual ownership data for the entire stock of privately owned rental units in France (the 2019 *FIDELI* database for c. 7 million rental dwellings used as primary residences.

Panels (a) and (b) of figure 1-4 illustrate this fact, mapping the residence distribution of individual owners of rental dwellings located in two large French commuting zones, Strasbourg and Toulouse.[23] While owners from populated and wealthy commuting zones like Paris or Lyon own a large number of units in both destinations, the striking pattern is the presence of substantial local bias, with the share of owners living in the commuting zone

[22]This regularity is reminiscent of the international finance literature, where foreign direct investment and portfolio holdings have been found to increase with measures of bilateral country proximity (Portes, Rey, and Oh, 2001; Portes and Rey, 2005). The presence of within-country home bias in investment also accords with the findings of Coval and Moskowitz (1999) for institutional investors stockholdings.

[23]The number of rentals owned in j for owners residing in i is measured from exhaustive cadaster information in the *FIDELI* database (see appendix 1.C), aggregated at the level of mainland France's 304 commuting zones ("*zones d'emploi*").

itself above 60 percent, and a clear geographic decay of portfolio shares with the distance between investor residence i and destination j. This more systematic finding is documented for close to 7 million private rentals in France in figure 1-5: the (log) number of units owned by investors from i in j decreases strongly with the distance between i and j.[24] Table 1.B.1 presents "gravity" estimates of the distance effect on rental property ownership, conditional on origin and destination fixed effects. The distance coefficient for predicting relative ownership probabilities hovers around unity; it is significant and economically substantial, with the effect largely mediated by social proximity. Figure 1.A.1 visualizes home bias for alternative distance measures (haversine distance, bilateral road distance) and aggregation levels (provinces, CZ).[25]

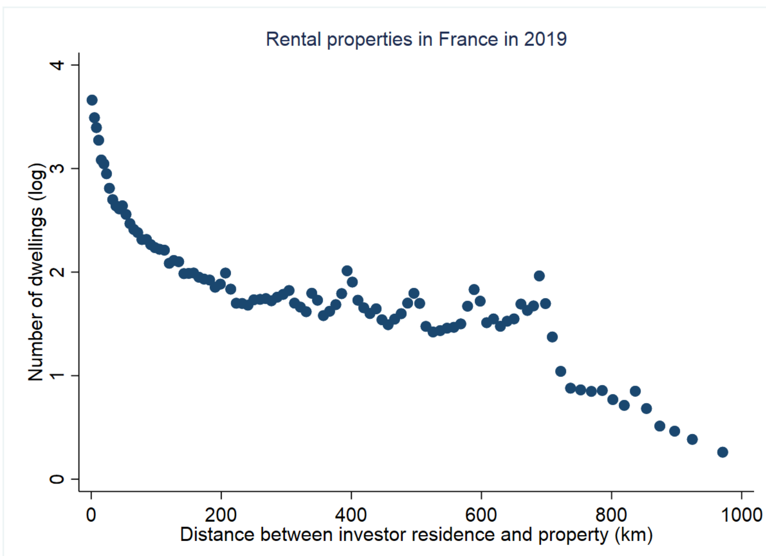

Figure 1-5: **Distance effects on ownership of rental housing**

Figure 1-5 documents the spatial concentration of housing ownership, using individual ownership data for the entire stock of privately owned rental units in France. It plots a binned scatter plot of the (log) number of units in municipality j owned by households residing in a municipality i, against bins of geographic haversine distance (in km, plotted up to a maximum distance of 1,000km) between the centroids of municipalities i and j. The municipality of residence of the owner is imputed from the municipality of the address to which the property tax (remitted by the owner) is sent.

[24] I compute the haversine distance between the town in which a property is located, and the town of the address to which the property tax is sent (a proxy for the residence of the owner, since owners are legally liable for property tax payments in rental units), for all private rental units in the country.

[25] Social connections are proxied by the Facebook friendship-based Social Connectedness Index (Bailey et al., 2018), only available across NUTS2-level provinces (*departements*) in France.

The over-representation of local ownership in the rental market is even more puzzling than international equity home bias. Most landlords are owner-occupiers themselves: they could better insure against local price risk by investing in distant locations with low correlation to domestic house prices.[26] Several factors could account for the strong observed home bias. Local landlords may have informational advantages relative to out-of-town buyers, through better timing, targeting, or negotiating strategies (Chinco and Mayer, 2016); they could save on intermediation and maintenance costs; or they could be better able to monitor tenants and alleviate the moral hazard inherent to the provision of rentals.[27] Without taking a stand on the underlying explanation, the large effect of distance on ownership suggests an imperfect mobility of rental capital supply across space: frictions reduce the subjective returns of landlords when investing in homes away from home.

1.3 Home-biased landlords: a spatial model of rental housing

The previous section provided stylized evidence that landlords, who are older and wealthier than the rest of the population, are mostly undiversified individual operators of rentals close to their own residence, facing significant capital market frictions. I outline here how an illustrative conceptual framework combining a standard location choice with the frictional portfolio decisions of home-biased landlords can rationalize the steady-state regularities documented in section 1.2, and deliver specific predictions for the quasi-experiment analyzed in sections 1.4 and 1.5 of this paper.

[26] Home-ownership itself may provide a long-term hedge against volatile local rents, for owners endowed with location-specific human capital (Sinai and Souleles, 2005). However, conditional on hedging rent risk in their own preferred location via owner-occupied housing, the ownership of *additional* rental assets in or near the same city is less straightforward to rationalize under perfect capital mobility.

[27] An alternative explanation involves long-term persistence. If movers lease, rather than sell, their past owner-occupied homes, and if mobility is limited across space, then the ownership structure of rental assets would exhibit local bias. Similarly, with bequests and low inter-generational spatial mobility, landlords will tend to live close to their parents' former homes, which they inherited and may now lease; and if landlords invest in local assets with a view to donating them to their own children, limited inter-generational spatial mobility and altruistic motives will again entail negative distance effects on rental property ownership. However, for landlords to prefer continued ownership of such rental assets, they must still exhibit some objective or subjective comparative advantage in their operation.

1.3.1 Setup: renters, savers, and developers

Environment Locations are indexed by $j = 1, \ldots, J$. A number $\bar{L}$ of agents are born at time t, and live for two periods in overlapping generations. Each generation is split into two fixed-types subsets: mobile hand-to-mouth worker-renters W, and immobile owners-savers O who can become landlords. Wages y_j^k (for $k \in O, W$[28]) and amenities B_j in each location are exogenous, while home prices, rents, and population are the endogenous outcome of mobility decisions, housing production, and portfolio allocation choices.

Workers-renters A share λ of households rent housing and work in final good production. Worker-renter behavior follows standard spatial equilibrium models (Redding and Rossi-Hansberg, 2017). These agents, indexed by ω, choose where to live at birth for the duration of their existence, earning the city-type specific wage income each period. Their utility is given by composite consumption in each of the two periods of their lives:

$$U_j^W(\omega) = (1-\beta^R)\ln(C_{t,j}^W \eta_j(\omega)) + \beta_R \ln(C_{t+1,j}^W \eta_j(\omega))$$

Renters are hand-to-mouth: they consume their entire wage income each period, and, with Cobb-Douglas preferences over housing and the (numeraire) final good, they spend a share γ of their income on rent. With constant local wages, consumption is the same each period and lifetime utility equals $\ln(C_{t,j}^W \eta_j(\omega))$. Workers elect a preferred location depending on wages for their type y_j^W, housing rent R_j, and idiosyncratic preference shocks $\eta_j(\omega)$, drawn from an extreme-value distribution with shape parameter ν and scale B_j denoting average amenities, as in Redding (2016). Standard arguments imply that the share of renters choosing to reside in j is given by:

$$\pi_j^W = \frac{B_j (y_j^W R_j^{-\gamma})^\nu}{\sum_k B_k (y_k^W R_k^{-\gamma})^\nu} \tag{1.1}$$

[28] Production in the model is linear in each type of local labor, so that type-specific earnings are constant.

In turn, the quantity of rental housing services required by a generation of workers-renters in location j is:[29]

$$H_j^{WR} = \frac{\pi_j^W \times \lambda \bar{L} \times \gamma y_j^W}{R_j} \tag{1.2}$$

Landlords-savers The remaining $(1-\lambda)\bar{L}$ agents are fixed-location owners-occupiers. An exogenous share $\bar{\pi}_i^O = \frac{L_i^O}{(1-\lambda)\bar{L}} > 0$ of them lives in each location i. To account for their fixed location, one may assume these agents are endowed at birth with unalienable owner-occupied housing in a given city - for example, through bequests from the last generation of owners. Unlike workers-renters, they can – and, in equilibrium, must – save. An owner-saver residing in i, indexed by z, earns a city-type-specific income y_i^O in period t ("working life"), and lives on the proceeds of his investments in $t+1$ ("retirement"). The utility of savers is increasing in consumption in both periods, with discount factor β_i^O:

$$U_i^O(z) = (1-\beta_i^O)\log(C_{t,i}^O(z)) + \beta_i^O \log(C_{t+1,i}^O(z))$$

Log utility implies each saves a constant fraction β_i^O of their first-period income y_i^O.[30] During their working life, savers face two potential subsets of assets. First, they can invest in "stocks" (representing the global financial market), receiving a baseline gross return r_K. Alternatively, they can purchase housing units in destination city j before leasing them to next-period renters. If they choose to do so, they pay the unit price P_j^H, before receiving (potentially subsidized) rents of $R_j S_j$ in the next period,[31] yielding a return $r_j^H = \frac{R_j S_j}{P_j^H}$.
Each constrained investor allocates their savings to their highest-return investment opportunity.[32] They are, however, subject to frictions, so that individual subjective returns from

[29]Because I focus on steady-state implications of the model, I mostly omit time subscripts where it is innocuous.

[30]In the absence of risk, this functional form choice for period felicity delivers constant savings shares independently of local expected returns, since income and substitution effects of varying asset returns compensate each other exactly.

[31]In the natural experiment of in the next sections, investment tax credits in j lower the purchase price for investors by $(1-s_j)$, but require rents to remain below a ceiling $\bar{R}_j$. The combination of these tools (a subsidy to the purchase price and a reduction in rents) in this simple model is isomorphic to a proportional subsidy S_j on rents received by landlords. Subsidies are funded with lump-sum nationwide taxes that do not distort worker location or saver investment decisions.

[32]To make the implications of the framework stark, and consistent with motivating facts on landlord wealth documented earlier, I assume savers are unable to borrow when taking advantage of high-return opportunities, so that their investment decision is constrained by their own savings – but I conjecture that any binding

an investment opportunity differ from market-level observed gross return rates.[33] Closely following Pellegrino, Spolaore, and Wacziarg (2021), the subjective returns to an investor z, living in i, from investing in an asset ℓ (a rental dwelling in location j delivering $r^{\ell} = r_j^H$, or a stock with $r^{\ell} = r_K$), are altered by an "iceberg" friction:

$$r_{z,\ell}(i,j) = r^{\ell} e^{-\lambda(z,\ell,i,j)}$$

where:

$$\lambda(z,\ell,i,j) = \delta \times \mathbb{1}(\ell \in H) \times d_{ij} - \epsilon_{z,l}$$

In the housing sector, investment frictions involve a systematic component increasing with the bilateral distance d_{ij} between the origin and destination locations. Moreover, all returns are affected by an idiosyncratic shock $\epsilon(z,\ell)$, at the saver z-asset ℓ pair level, drawn from a nested Frechet distribution.[34] Since each saver allocates their entire savings to their highest return opportunity, this delivers a tractable expression for w_{ij}, the share of investors from i who purchase residential assets in j. This proportion is also, by the law of large numbers,

contractual frictions limiting their ability to access external capital to a multiple of their own private savings (as would result from standard repayment monitoring constraints a la Holmstrom and Tirole (1997)) would have similar qualitative consequences.

[33] In practice, in the case of the housing market, subjective return variation from a housing investment project in a given destination city may correspond to any investor- and asset-specific idiosyncratic payoffs. As mentioned in section 1.2, these could stand for an individually-varying hassle cost to examine, purchase or maintain the property. Alternatively, akin to the differentiated liquidity benefits assumed in the literature on "convenience yields" for financial securities (Feenstra, 1986; Krishnamurthy and Vissing-Jorgensen, 2012), they could represent individual variation in the perceived probability of enjoying a given asset as a secondary residence, or of bequeathing it to one's children. For other outside assets ("stocks"), any variation in individual-asset specific transaction costs or individual preferences could generate the type of return dispersion assumed here. Absent any financial frictions, only investors with the highest subjective return would invest in a given asset, borrowing from all other savers, an implication that runs counter to observed patterns of portfolio allocation.

[34] Denoting each sector (housing and outside assets) by s, and N_s the number of potential assets in each sector, the distribution of shocks satisfies:

$$F\left(\epsilon_{z,\ell}(s,k)\right) = \exp\left[-\sum_{s\in(K,H)}\left(\sum_{k=1}^{N_s}(\epsilon_{z,l}(s,k))^{-\theta}\right)^{\frac{\phi}{\theta}}\right]$$

This type of nested shock distribution, with individual extreme-value correlated preference shocks giving rise to a nested logit choice probability (Verboven, 1996), has been used to model preferences across goods and varieties (Atkeson and Burstein, 2008), firms' funding choices (Herreno, 2020), or workers' employment location decisions (Berger, Herkenhoff, and Mongey, forthcoming). The (inverse) dispersion of shocks *across* the housing and non-housing nests is parameterized by ϕ; while the (inverse) dispersion of return shocks *within* housing assets depends on θ.

the share of wealth in i ($W_i = (1-\lambda)\bar{L}\pi_i^O \beta_i^O y_i^O$) invested in j housing:

$$w_{ij} = \underbrace{\frac{\tilde{m}_{Hi}^{\phi}}{\tilde{m}_{Hi}^{\phi} + r_K^{\phi}}}_{\text{Housing share } HS_i \text{ in } i \text{ assets}} \times \underbrace{\frac{\left(e^{-\delta d_{ij}} r_j^H\right)^{\theta}}{\sum_{j'} \left(e^{-\delta d_{ij'}} r_{j'}^H\right)^{\theta}}}_{\text{Share } s_{ij} \text{ of } j \text{ in } i\text{'s housing portfolio}} \tag{1.3}$$

where $\tilde{m}_{Hi} = (\sum_{j'} \left(e^{-\delta d_{ij'}} r_{j'}^H\right)^{\theta})^{\frac{1}{\theta}}$ is a theory-consistent index of net housing returns for investors living in i. The portfolio weight w_{ij} is the product of two terms. The first is the overall probability HS_i of investing in housing (rather than stocks) for investors from i; the second is the probability (for i savers) of buying property in j, conditional on investing in any housing asset. A higher θ makes housing investment more elastic to returns across space, while a higher ϕ makes housing and non-housing investments more substitutable from the point of view of investors. The demand for housing assets in location j at time t (which is the time-$t+1$ supply of rental services) is given by the aggregation of investors' purchases from each location i ($H_{i\to j}$):

$$P_{jt}^H H_{jt} = P_{jt}^H H_{j,t+1}^{SR} = \sum_i P_j^H H_{i\to j} = \sum_i \pi_i^O (1-\lambda)\bar{L} \times \beta_i^O y_i^O \times w_{ij} \tag{1.4}$$

Housing production by developers Housing used for rental services is produced by competitive developers using local land T_j, as well as materials X_j (units of the numeraire). The price of local land is r_j^T, with a limited supply assumed to depend on exogenous characteristics $\bar{T}_j$ and its own price: $T_j = \bar{T}_j (r_j^T)^{\sigma_j}$. The production function for new housing is Cobb-Douglas, with land share ξ_j: $H_j = (\frac{T}{\xi_j})^{\xi_j} (\frac{X}{1-\xi_j})^{1-\xi_j}$. Solving for the maximization problem of competitive developers, the land price is given by $r_j^T = (P_j^H)^{\frac{1}{\xi_j}}$, and the supply of housing assets can be shown to follow[35]:

$$H_j = \frac{\bar{T}_j}{\xi_j} (P_j^H)^{\frac{1+\sigma_j - \xi_j}{\xi_j}} \tag{1.5}$$

[35] Equilibrium land prices are log-linear in local housing prices since materials are the numeraire; and the output price elasticity of housing production is increasing in the land conversion responsiveness, and decreasing in the land share of construction.

Definition: stationary equilibrium *Given wages y_j^W, y_j^O, amenities B_j, buildable land characteristics $\bar{T}_j$, subsidies S_j, and owner residence distribution π_j^O, a steady-state spatial equilibrium of this economy with constant population is an allocation of renters L_j^W, housing production H_j, housing prices P_j^H, and rents R_j such that, in each generation:*

- *The number of workers-renters living in each destination market j, L_j^W, is constant and satisfies their optimal location choice, summarized by π_j^W in equation 1.1.*
- *The savings allocated by investors from each origin city to each destination housing market are constant, and, given returns, satisfy the optimality condition for individual investors, summarized by the weights w_{ij} in equation 1.3.*
- *The demand for new housing assets at time t is given by investor asset allocation 1.4; the supply of new housing is determined by developer competition 1.5; and the market for new housing production and land both clear each period.*
- *The market for rental services clears each period, with the demand H_{jt}^{WR} given by renter aggregation 1.2 at time t, and the supply $H_{j,t}^{SR}$ given by portfolio allocation choices 1.4 at time $t-1$.*

Discussion of assumptions The framework outlined above is highly stylized, but the hypotheses made on renters and landlords are rooted in real-world regularities.

First, the model postulates that only renters move across space. Using Census data, appendix 1.E.1 provides a trove of empirical evidence at the individual and town level that renters are much more prone to moving than owner-occupiers, likely due to the presence of transaction and mobility costs. Tenure duration in the private rental sector is substantially lower than in both public rental and owner-occupied housing segments.

Second, the model assumes a perfect segmentation of the population between fixed-location, "older", landlords, and mobile, hand-to-mouth, "younger" renters. While clearly an extreme case, since many households may transition from renting to owning over the course of their life cycle, this distinction is qualitatively consistent with the descriptive stylized facts presented in section 1.2 and table 1.E.4 in appendix 1.E.2: individual landlords are substantially older and wealthier than renters – whose savings are almost nil.

Third, the framework distinguishes between the production of housing by developers, and the delivery of rental services by landlords. This structure is peculiarly suited to the analysis of the French (and Western European) rental market, where builders sell ownership rights for dwellings to individual investors before or immediately after construction. The United States has a larger proportion of corporate landlords who erect, own and operate entire multi-family buildings, but individual landlords still play a substantial role, representing 7 percent of all US households and operating 41 percent of all rental units.[36] The separation of competitive production, on the one hand, and lumpy ownership and operation, on the other, is also relevant in the US.

Fourth, instead of a standard portfolio choice under risk, spatial frictions to the operation of housing assets are assumed to reduce the private value of operating more distant rental investments, consistent with the home bias documented in section 1.2. This geographic segmentation also implies that many local owners of rental property may be infra-marginal in their portfolio allocation.

1.3.2 Model implications: the steady-state

The cross-section of cities The model matches key steady-state regularities of the rental market documented in section 1.2. By making assumptions on the frictional nature of housing investment across locations, it generates a cross-sectional spatial pattern of ownership and rent-to-price ratios consistent with the data.

Proposition 1.3.1. *(i) In equilibrium, the network for housing asset ownership exhibits a gravity structure, summarized by:* $H_{i \to j} = A_i \times \tilde{A}_j \times e^{-\delta\theta d_{ij}}$. *(ii) In a two-city symmetric version of the model, where two locations A and B only differ only by their share of saver wealth* $\pi^O_A y^O_A > \pi^O_B y^O_B$, *the city with a higher "investor portfolio access" (A) exhibits a lower rent-to-price ratio* $\frac{R_A}{P^H_A} < \frac{R_B}{P^H_B}$, *a larger housing supply* ($H^W_A > H^W_B$) *and a higher population of workers-renters* ($L^W_A > L^W_B$). *(iii) These cross-sectional differences vanish in the absence of distance frictions* ($\delta \to 0$).

- *Gravity in housing investment* Because of the dispersion of owners' idiosyncratic preferences, and the presence of distance frictions in investment, the model delivers a gravity

[36] These summary data for the rental market in the United States are computed from the Census's Rental Housing Finance Survey for 2018 and the IRS's Statistics of Income for the same year.

structure of the rental housing ownership network similar to the one observed in the third fact of section 1.2 and quantified in table 1.B.1. Specifically, housing in j owned by investors from i has the following gravity structure[37]:

$$H_{i\to j} = \kappa \times \underbrace{\pi_i^O(1-\lambda)\bar{L} \times \beta_i^O y_i^O \frac{\tilde{m}_{Hi}^{\phi-\theta}}{\tilde{m}_{Hi}^{\phi} + r_K^{\phi}}}_{\text{Origin FE: } A_i} \times \underbrace{\frac{(H_j)^{-\theta\left(\frac{\xi_j}{1+\sigma_j-\xi_j}+\frac{1}{1+\gamma\nu}\right)}\left(B_j (y_j^W)^{1+\nu}\right)^{\frac{\theta}{1+\gamma\nu}}}{\left(\frac{\xi_j}{T_j}\right)^{\frac{\xi_j\theta}{1+\sigma_j-\xi_j}}}}_{\text{Destination FE: } \tilde{A}_j} \times \underbrace{e^{-\delta\theta d_{ij}}}_{\text{Bilateral friction}}$$

- *Decreasing rent-price ratios with proximity to wealth* Locations in close proximity to savers feature a "savings glut" in the housing sector, with a large supply of capital and low returns.[38] While the proposition is stated for the symmetric case of two cities, the logic is more general. More numerous (high π_i^O) or wealthier (high y_i^O) savers in locations i "close" to j (i.e. with low effective bilateral frictions d_{ij}) supply relatively more potential housing capital to j. All else equal, this improved "investor portfolio access" (summarized by a higher $\sum_i e^{-\delta d_{ij}} \beta_j^O \pi_j^O y_j^O$) lowers the equilibrium return to housing in j.[39]

- *Poor people in wealthy cities* Spatial equilibrium models often ponder how to account for the observed presence of low-income workers in expensive agglomerations, since competition for land may be expected to induce full segregation of types.[40] With segmented capital markets, a larger number of affluent savers directly lowers the rent-price ratio, a pecuniary cross-agglomeration force pushing for the in-migration of low-income renters – conditional on other features of housing production. Locations with a larger supply of rental housing (as a consequence of their proximity to investors)

[37] κ is a composite constant dependent on model parameters.

[38] This relates the findings in my paper to recent evidence in Kuchler et al. (forthcoming) on how higher social connections to locations replete with institutional investors tend to increase firm valuation. This "investor portfolio access" effect is also similar to how, in the presence of trade in goods, locations where transportation infrastructure provides a high market access to efficient producers feature lower effective consumer price indices (Donaldson and Hornbeck, 2016).

[39] Separately, in the many-cities version, the net effect of investor portfolio access on rents and on prices respectively is ambiguous, as it depends on the respective elasticities of land supply σ_j and location choices ν, along with the investment reallocation parameters θ and ϕ.

[40] Asking this exact question, Glaeser, Kahn, and Rappaport (2008) provide an alternative answer focusing on the role of access to public transportation.

host more mobile (lower-income and younger) workers, consistent with the evidence documented in appendix 1.E.1.[41]

Graphical intuition for proposition 1.3.1 One can visualize the steady-state of this economy as the combination of a real or "physical" spatial equilibrium schedule, in which workers make location choices as a function of rents and developers supply housing units as a function of prices; and a "financial" optimality condition, in which savers allocate their wealth to housing assets depending on their returns (net of subsidies and frictions).

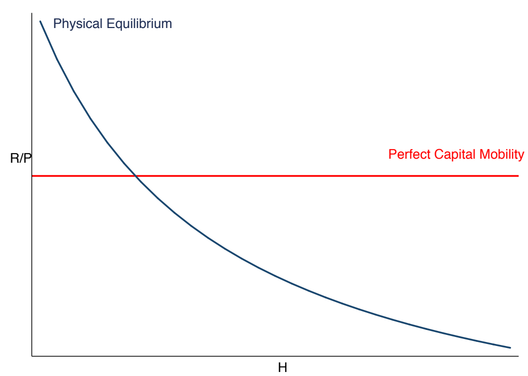

(a) Perfectly elastic capital supply

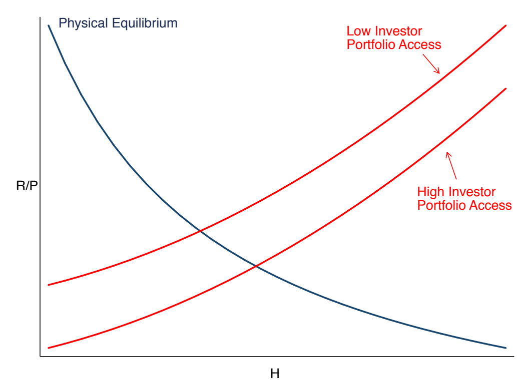

(b) Imperfectly elastic capital supply

Figure 1-6: **Rent-price ratios and rental supply: steady-state**

Figure 1-6 summarizes the role of imperfectly elastic capital supply for the equilibrium supply of rentals (on the x-axis) and rent-price ratio (plotted on the y-axis). In both panels, the decreasing schedule (blue line) summarizes the "physical" equilibrium (a higher quantity of rentals leads to lower rents and higher prices, and thus to a lower rent-price ratio). The capital supply schedule is plotted (in red) for two cases. Panel (a) presents a generic city in a financially integrated country with an exogenous and constant required return for landlords. Panel (b) plots the case of an imperfectly elastic supply of capital, and shows the impact of spatial frictions on steady-state returns and supplies of rentals for two locations with (resp.) a high (low) equilibrium access to investors' portfolios.

On the "physical" side, all else equal, a larger quantity of rental housing in a location requires *lower* rents to clear the market for rental services ("going down" the demand curve

[41] In an extended version of the model with mobile landlords, a dynamic counterpart of this prediction would exist. If landlords were to also move in the presence of amenity differentials, their share in j, π_j^O, would rise in response to an increase in local amenities B_j. Assuming such an amenity shock initially draws both landlords and workers to j, savers now face lower frictions to rental investment in j on average. This pushes the rent-to-price ratio lower in j (and higher in locations experiencing reduced investor portfolio access), dampening the direct impact of the population inflow on rental costs in the city. This force may partially account for the long-run drop in rent-price ratios specific to "superstar cities", where prices have risen faster than rents (Gyourko, Mayer, and Sinai, 2013; Hilber, Mense, et al., 2021).

of workers for rentals), and *higher* unit prices to clear the market for housing production ("going up" the supply curve of developers for new units). Combining these, the "physical" equilibrium delineates a *decreasing* schedule between local housing quantities and the rent-to-price ratio.

In the absence of financial frictions, the intersection of this schedule with an exogenously defined required return r^K for absentee landlords would equalize the rent-to-price ratio in every location. With imperfect capital mobility, however, the "financial" side delivers a location-specific upwards-sloping supply of rental capital in the local rent-price ratio: attracting further investment to a city (by drawing in more distant investors and reallocating non-housing capital) requires higher gross returns. As a consequence of this origin-destination-specific capital supply, locations with a higher aggregated "investor portfolio access", a measure of their proximity to investors, feature a capital supply curve shifted outwards, and lower returns to rental investment. Figure 1-6 summarizes this discussion: the left panel describes the case of integrated capital markets for a generic city, while the right panel shows local quantities and rent-price ratios for two cities, with a low or high access to investors' portfolios, under imperfect capital mobility.

1.3.3 Model implications: the effect of subsidies

Response to landlord subsidies The model generates direct predictions for the consequences of raising S_j, the subsidy to rental investment in location j, to which I turn in the quasi-experimental evidence obtained in sections 1.4 and 1.5. In particular, the theory predicts not only that subsidies raise investment and housing supply in the targeted location, but also that their effects operate through a rise in long-distance investor involvement, and are stronger in locations exhibiting more openness to outside investors.

Proposition 1.3.2. *(i) An increase in the subsidy $dS_j > 0$ to the provision of rental services in location j raises housing production H_j, renter population L_j^W, and the price of both land (r_j^T) and new dwellings (P_j^H) in j; it also increases net-of-tax returns to investment $r_j^H = \frac{R_j S_j}{P_j^H}$.*

(ii) For $\phi \frac{r_K^\phi}{\tilde{m}_{Hi}^\phi + r_K^\phi} < \theta$, the partial equilibrium effect of the subsidy $\frac{\partial \log P_j^H H_{i\to j}}{\partial \log S_j} = s_{ij}\phi(1 - \mathrm{HS}_i) + (1 - s_{ij})\theta$ is stronger *for investment flows $H_{i\to j}$ stemming from origin locations i with a* lower *initial ownership share s_{ij}.*

(iii) The response of overall local housing provision to subsidies $\frac{\partial \log P_j^H H_j}{\partial \log S_j}$ *is* larger *in treated locations with a* lower *share of self-ownership* s_{jj}.

- *Effects of investment tax credits on investment, housing production, prices, and population* A location-specific subsidy S_j for the purchase of rental property in j shifts the allocation of savings towards housing in general, and location j dwellings in particular. The subsidy raises housing production H_j and sales $P_j^H H_j$ in j, leading to higher land (r_j^T) and new unit prices (P_j^H). The induced increase in the housing stock triggers inwards renter mobility, inducing population growth $dL_j^W > 0$ in treated locations.

- *Incomplete capitalization* The usual view of exogenously determined post-tax returns leaves no room for any incidence of place-based policies on savers. Under an imperfectly elastic capital supply, however, the subsidy is not fully capitalized in rents and prices. Post-tax returns $\frac{R_j S_j}{P_j^H}$ in treated locations rise to induce new investors to switch to rental investment in j at the extensive margin. Therefore, the pre-tax return $\frac{R_j}{P_j^H}$ drops by less than the full amount of the subsidy (either through higher acquisition prices or lower rents), and part of the incidence of place-based policies falls on infra-marginal landlords who would have invested in j even absent the subsidy and enjoy higher net returns as a consequence of it. Figure 1-7 evidences this phenomenon, showing how the outwards shift in the supply of capital fails to be fully capitalized in pre-tax returns under an upwards-sloping supply of capital.[42]

[42] The full capitalization under perfect capital mobility is reminiscent of the effect of investment tax credits on prices in House and Shapiro (2008). Due to inter-temporal substitution in investment, the ability to substitute capital goods acquisitions across *time* implies that firms' demand for long-lived assets is almost infinitely elastic to temporary changes in the tax treatment of investment, and the price of equipment goods should reflect the full value of tax incentives. Similarly, under the (violated) null assumption of perfect capital mobility across *space*, the ability of investors to perfectly arbitrage return differentials between subsidized and untreated locations would imply a demand for residential capital assets that is perfectly elastic to returns across places, and, in turn, full capitalization of the value of place-based incentives in prices.

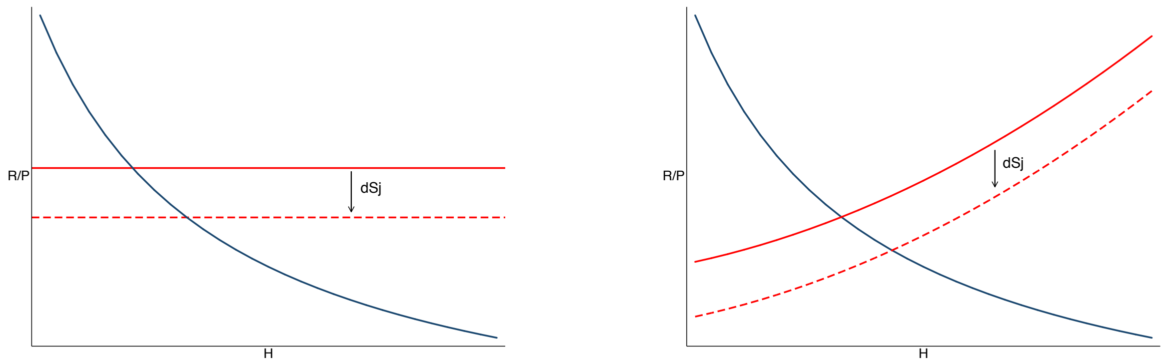

(a) Perfect capital mobility

(b) Imperfect capital mobility

Figure 1-7: **Imperfectly elastic capital supply and the effects of landlord subsidies**

Figure 1-7 summarizes the role of imperfectly elastic capital supply for the response of housing investment and returns to landlord subsidies. Panel (a) presents a generic city in a financially integrated country: returns fully capitalize the value of the landlord incentive, leaving post-tax returns unaffected, either through a rise in prices or through lower rents. Panel (b) plots the case of an imperfectly elastic supply of capital: here, while the production of new housing also rises, part of the incidence of subsidies falls upon infra-marginal savers who receive higher post-tax returns.

- *Rise in long-distance investment* An implication of equation 1.3 is that the partial equilibrium elasticity of housing investment flows from i to j varies systematically with the portfolio share of j in i, s_{ij}:

$$\frac{\partial \log P_j^H H_{i\to j}}{\partial \log S_j} = s_{ij}\phi(1-\mathrm{HS}_i) + (1-s_{ij})\theta$$

When the conditional share $s_{ij} = \frac{\left(e^{-\delta d_{ij}} r_j^H\right)^\theta}{\sum_{j'}\left(e^{-\delta d_{ij'}} r_{j'}^H\right)^\theta}$ of j in i's portfolio is high (i.e. when bilateral frictions d_{ij} are low), the partial equilibrium elasticity of investment to local returns is mostly governed by the reallocation from outside assets to housing, which depends on ϕ. On the contrary, when s_{ij} is low (more distant locations), the partial equilibrium responsiveness of investment flows to returns mainly depends on θ, the parameter that affects the ease of "within-housing", lower-nest reallocation. Intuitively, nearby investors are mostly infra-marginal: most choose to invest in j independently of the subsidy, and their response, if any, operates through the difficult reallocation of non-housing to housing assets. On the other hand, for long-distance investors, a shock to returns in j shifts their large non-j *housing* investments towards the treated location.[43] With $\theta > \phi(1 - \text{Housing share}_i)$, a higher subsidy thus raises the proportion of long-distance investors.

- *Heterogeneous responses by access to outside capital* Aggregating flows from all locations, the response of housing provision to a place-based subsidy is heterogeneous across space. Since the response is stronger for longer-distance investors as described above, the locations that benefit most from the subsidy are those with more dispersed initial ownership shares, where a larger baseline proportion of housing is owned by these more distant landlords.[44] Formally, this openness to outside capital is summarized by a model-implied measure of the dispersion of ownership: $1 - \sum_i \frac{H_{i \to j}}{H_j} s_{ij}$, the sum of i investors' housing portfolio shares in j, weighted by the baseline share of investors from i in j.[45]

1.4 Quasi-experimental evidence: setup, strategy, and data

To assess the causal impact of tax subsidies on the spatial allocation of residential investment, and show that they are consistent with the mechanism described in section 1.3, I turn to quasi-experimental evidence on landlords' decisions. I rely on exogenous place-specific changes in eligibility to a tax credit for buy-to-let investment, and estimate its consequences for the housing market, landlords, and local communities. This section describes the institutional setup I exploit for causal inference (1.4.1), defines my empirical strategy (1.4.2), and summarizes the data (1.4.3).[46]

[43] This phenomenon is similar to the findings of Berger, Herkenhoff, and Mongey (forthcoming), who study the response of employment to firm wages. When a firm has a large employment share in a given market, most of the response to higher wages comes from cross-market employment reallocation. Here, when housing in j represents a large portfolio share in i, the bilateral investment response to return shocks is drawn from cross-asset (rather than "within-housing") reallocation.

[44] As in standard spatial frameworks, the heterogeneous response to location-based subsidies also depends on "physical" supply conditions arising from locally varying physical land abundance and intensity in the production function (denoted by the composite elasticity $\frac{1+\sigma_j-\xi_j}{\xi_j}$).

[45] This measure is highly correlated with $1 - s_{jj}$, the complement of the own-ownership share in the housing stock. It is reminiscent of the heterogeneity in local employment elasticities in Monte, Redding, and Rossi-Hansberg (2018) that stems from variation in the "own-employment" share and makes cities more open to commuting more able to accommodate labor demand shocks.

[46] Appendix 1.C includes additional information on the data sources used in the paper. Appendix 1.D provides more detail on the institutional setup and French housing supply incentives.

1.4.1 Institutional background

The *Pinel* housing tax credit In January 2013, French Housing Minister Cecile Duflot introduced a conditional individual income tax rebate for taxpayers purchasing newly built real estate in eligible locations, if they offered affordable rents for nine years.[47] The *Duflot* law was revamped in September 2014 by Duflot's successor, Sylvia Pinel. Under the new *Pinel* tax credit, individual landlords purchasing a new dwelling receive a personal income tax reduction of up to 21% of their investment if they lease the unit for a minimum of six years below a regulatory rent ceiling, and tenant income does not exceed a threshold amount.[48] Rent and income ceilings both vary with the location of the municipality (see *infra*). Annual investment is limited to two dwelling units and a maximum of €300,000 (ca. USD 350,000) per year.[49] The reduction only applies up to a purchase price of €5,500/sq.m., creating a kink in the credit. Several changes made the *Pinel* scheme more favorable than the *Duflot* mechanism. A more flexible rental duration commitment (6, 9, or 12 years) was available under the *Pinel* tax credit (only 9 years with the *Duflot* scheme); a prohibition to rent the unit to one's relatives was removed; and a reshuffling of eligible areas (see *infra*) broadened access to the scheme. Table 1.B.3 provides a range for the net present value of the landlord incentive between 9 and 15 percent, accounting for varying interest rates and rent discount stringency estimates.

Place-based targeting Eligibility, maximum rents, and income ceilings under the *Pinel* scheme follow strict place-based regulations (known as the *zonage Pinel*) by municipality. Nationwide rules classify France's 35,000 municipalities into five classes, ranked by descending order of government-assessed housing market "tightness". These classes, illus-

[47] Relative to its predecessor, the *Scellier* scheme implemented in 2009, the *Duflot* law had lower levels of investment tax credits, and more stringent affordability requirements, with rent and income ceilings close to 20 percent lower (see appendix 1.D.1).

[48] The tax reduction is computed as 2% of the purchase price each year after the start of the lease, for up to six or nine years (depending on the taxpayer's initial affordable rent duration commitment), with an option for a 1% additional annual reduction for years 10 to 12, if the taxpayer extends her affordability commitment - yielding a maximum reduction of 21%.

[49] The maximum subsidy is therefore $0.21 \times 300,000 = 63,000$ euros over twelve years. The *Pinel* scheme is a non-refundable tax reduction. Because of the indivisibility of housing units documented in section 1.2, the 2 percent annual non-refundable reduction only fully benefits landlords with a substantial tax liability to offset: around 60% of beneficiaries are in the top decile of household income, and 90% in the top tercile, according to government reports (Inspection Generale des Finances, 2019).

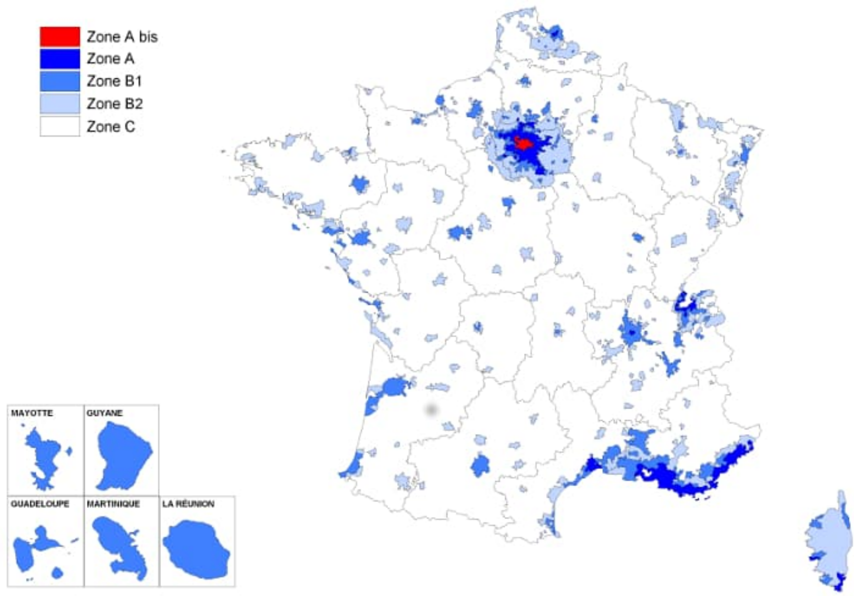

Figure 1-8: ***Pinel* zoning**

Municipalities were mapped into one of five areas (Abis, A, B1, B2, C), determining both their eligibility to the *Pinel* scheme, and rent guidance and income restrictions. Areas are ranked in descending order of estimated property market "imbalances". The classification was widely reshuffled with the implementation of the *Pinel* scheme in October 2014 Source: French Housing Secretariat.

trated in figure 1-8, range from *A bis* (Paris and its closest neighbors), to *A* (the broader Paris region, Marseille, Lyon, and some touristic towns on the French Riviera), *B1* (several other large cities, including Toulouse, Bordeaux or Strasbourg), *B2* (most medium-sized cities and their surroundings, like Dunkirk, Saint-Etienne, or Le Mans), and *C* (all 30,000 other towns). Under the scheme, each "zone" is subjected to different rent ceilings and maximum renter income. Rents per square meter are capped at $R = M \times C$, the product of a coefficient M (decreasing with dwelling size and capped at 1.2: $M = \min[\frac{19}{\text{Size}} + 0.7, 1.2]$), and a zone-specific ceiling amount C.[50] To accompany the roll-out of the *Pinel* scheme, a September 2014 government ruling modified the allocation of municipalities to the five classes, leading almost exclusively to zone upgrades and thus to less stringent rent and income ceilings for municipalities representing about 10 million inhabitants, close to a sixth of the total French population.[51] As shown in appendix 1.D.2, there is a wide heterogeneity in incomes and

[50] Appendix 1.D.1 provides additional details on the computation of the cap for local rents and incomes.

[51] 162 B1 municipalities (with close to 4 million residents, including the major cities of Lyon, Lille or Marseille) and 25 B2 municipalities (approx. 0.2 million res.) were upgraded to A. 8 municipalities (approx. 0.2 million res.) were upgraded from A to A bis. 358 B2 municipalities (hosting close to 3.4 million residents) and

rents among municipalities *within* a zone, and a significant overlap in housing market conditions and demographics *across* them. Nonetheless, applicable rent and income ceilings are uniform within a *Pinel* zone, and differ across them, yielding some variation in effective local stringency.[52]

Policy timing and reforms The *Duflot* tax rebate ran from January 2013 to August 2014, followed by the *Pinel* law after September 2014. Both policies were available, *de jure*, for the purchase of any new unit located in the A, A bis, and B1 areas. Units in C towns were not eligible to either scheme. Eligibility to the *Duflot* law in the approximately 3,900 B2 towns was phased out after June 2013. However, B2 towns could request an agreement from regional authorities in order to participate in the scheme beyond June 2013.[53] Authorizations were granted to around 900 B2 municipalities,[54] whose property markets were assessed to be closer to (always eligible) B1 municipalities. A purchase in an "agreed" B2 municipality (henceforth a "*B2+*" town) was thus eligible to the *Duflot* incentive from the date of the agreement until August 2014, and to the *Pinel* credit afterwards, while purchases in other B2 cities (henceforth a "*B2-*" town) were never eligible to the *Pinel* subsidy.

The *Pinel* scheme was initially due to expire on December 31st, 2016, but was extended in late 2016 until December 31st, 2017. In December 2017, the French parliament extended the scheme again until December 2021, but restricted it to areas A, A bis, and B1 to reduce its cost, thus excluding *B2+* municipalities.[55] To smooth the transition for construction projects already under way in *B2+* cities losing the benefit of the policy, sales in the *B2+* area remained

167 C municipalities (approx. 0.7 million res.) were upgraded to B1. 462 municipalities (approx. 2.4 million res.) were upgraded to B2. Conversely, only 38 municipalities (40.000 residents) were downgraded from A to B1, and 619 small municipalities (0.6 million res.) were downgraded from B1 to B2. No town was downgraded to C.

[52]Table 1.B.2 and 1.D.2 provide further summary information on the distribution of population, demographic characteristics, rents, and incomes across municipalities by zone and agreement status.

[53]The competent authority was the prefect, the State's representative in each of France's 13 regions.Prefects would grant authorizations based on their qualitative appreciation of "tightness" in the local housing market.

[54]Close to sixty percent (530) of these agreements were put in place in 2013; more than thirty-five percent were granted in 2014 and 2015, with the remaining agreements given after 2016. 20 (out of a bit more than 30,000) municipalities in the C area also obtained a special agreement. A handful of municipalities gained eligibility when a non-eligible town merged with an eligible one.

[55]The costs of tax breaks for landlords or corporate developers makes them a regular target of budget-balancing proposals, both in Europe and in the United States: the Congressional Budget Office regularly mentions "Repeal the Low-Income Housing Tax Credit" (link) as an effective revenue-enhancing avenue in its annual *Options for reducing the deficit*.

eligible to the tax break until December, 31st 2018 (later extended for technical reasons to March 15th, 2019), as long as the building permit had been requested before December 31st, 2017.[56]

1.4.2 Empirical strategy

Two natural experiments: entry and exit

Pinel scheme subsidies varied over time with a location's eligibility to the policy. Nationally mandated changes in coverage by the policy generate sharp variation in the tax treatment of residential investment, plausibly unrelated to local demand conditions or government priorities – a place-based setting that differs from a large literature focused on the role of tax or subsidy competition among communities deliberately offering different location incentives (Slattery and Zidar, 2020).[57] However, eligible A, A bis, B1 and B2+ municipalities were characterized by distinct housing prices and overall market trends, than untreated B2- and C towns, since the classification was constructed to correspond to relative property market tension (as shown by the starkly different characteristics of each group summarized in appendix 1.D.2 and table 1.B.2). Comparing the evolution of outcomes for all eligible and all ineligible towns would require strong parallel trends assumptions unlikely to be met, and would not disentangle the effect of subsidies from substantial composition differences in local dynamics. Fortunately, the government-designed partitioning of French municipalities into areas with comparable housing market trends builds natural control groups, for which differential dynamics in outcomes over time can credibly be attributed only to diverging subsidy eligibility. Specifically, I exploit two natural experiments linked to the timing and

[56] At the same time, the policy was expanded, starting January 2018, to cover the so-called CRSD (*Contrats de Redynamisation des Sites de Defense*) towns hosting former defence bases and receiving transitional financial support from the government when military installations closed. Therefore, a few B2+ municipalities with an agreement *remained* eligible to the policy, a feature I take into account when estimating the effect of the removal of the policy. Other B2- or C towns which had not hitherto been eligible *gained* access to it. I assess the effect of the subsidy in formerly ineligible (B2-/C) CRSD towns in appendix 1.D.3.

[57] The open-ended nature of the French buy-to-let tax credits, conditional on zone eligibility, is reminiscent of the Opportunity Zones implemented in the United States by the TCJA in 2017. This departs from schemes predicated on fixed local allotments of tax credits, such as the LIHTC. Incentives embedded in the LIHTC have been hard to estimate empirically, since a fixed amount of credits allocated to states are awarded as grants to projects at the discretion of local authorities. Quasi-experimental approaches to estimate the investment response have relied on a cross-sectional discontinuity in the LIHTC formula (Baum-Snow and Marion, 2009), which depends on the proportion of low-income households in a Census tract, and defines Qualified Census Tracts and Difficult Development Areas.

place-targeting of the policy:

1. **Loss of eligibility (*exit-DiD*)** To assess the short-run effect of tax credits on buy-to-let investment and eligible home sales and prices, I use a difference-in-differences strategy exploiting the *end of eligibility* for B2+ towns in December 2017, comparing formerly treated B2+ towns to B1 municipalities that remained eligible after the policy was scaled down. The exogenous differential loss of eligibility allows for a direct estimate of the response of residential investment by individual landlords. I label this difference-in-differences around the timing of eligibility loss the "*exit-DiD*" approach.

2. **Entry into eligibility (*entry-DiD*)** To explore medium-term consequences of tax subsidies for the broader local real estate market, housing conditions, and local communities, I estimate a difference-in-differences model around the timing of *entry* at the start of the *Pinel* scheme, leveraging the roll-out of agreements to B2+ municipalities, and comparing them only to a large permanent control group of never-treated "B2-" municipalities. I label this study of entry into policy eligibility the "*entry-DiD*" approach.

Preferred specifications

Difference-in-differences around loss of eligibility I estimate the short-term effect of removing access to the *Pinel* tax credit, using the *exit-DiD* approach. In this specification, I take advantage of the exogenous loss of eligibility in B2+ towns decided by the national government to directly estimate two outcomes of interest: the residential investment response, and the price effects of losing the subsidy. Since municipalities in the B1 area remain treated after the 2017 reform, while B2+ towns were no longer eligible, the *exit-DiD* specification restricts the sample to B2+ and B1 municipalities only. For an outcome y_{ct} in municipality c and time period t, I estimate a full leads-and-lags regression in the window $[\underline{T}, \overline{T}]$ leading to and following the end of eligibility:

$$y_{ct} = \alpha_c + \gamma_t + \sum_{\substack{k=\underline{T} \\ k\neq -1}}^{\overline{T}} \beta_k \mathbb{1}_{c\in \text{B2+}} \times \mathbb{1}_{t=k} + \Gamma X_{ct} + \epsilon_{ct} \tag{1.6}$$

where γ_t indexes a period fixed effect accounting for common time trends, α_c is a municipality fixed-effect that averages out municipality-level, time-invariant determinants, X_{ct} is a vector of time-varying controls at the municipality level (including initial municipality characteristics interacted with time fixed-effects), and the coefficients β_k measures the time-varying impact of eligibility removal on outcomes. The end of eligibility is the date after which units in B2+ locations were no longer eligible – corresponding to early 2019 for sales, and to December 2017 for permit requests. The key identification assumption is that, conditional on other characteristics, B2+ municipalities would have followed parallel trends to *always-eligible* B1 cities after that time, were it not for losing *Pinel* eligibility, since B2+ cities were deemed by regional prefects to more closely fit into the B1 classification. The leads-lag specification allows me to graphically visualize any differential pre-trends at high frequency *before* the removal of eligibility.

Difference-in-differences around entry After providing evidence on the immediate impact of losing eligibility to the policy for B2+ cities, I turn to its longer-run effects by exploiting the differential *entry* of B2+ cities – relative to B2- – after October 2014. I first confirm the quantitatively similar effect on sales at the intensive and extensive margin, before showing how the local outcomes for long-distance investment, the housing stock, and renter mobility track entry into eligibility among B2+ cities receiving an agreement, relative to a permanent control group of never-treated B2- locations. In this *entry-DiD* approach, I restrict the sample to all B2 municipalities in the period lasting until the end of eligibility, and use the following specification:

$$y_{ct} = \alpha_c + \gamma_t + \beta_R \mathbb{1}_{c \in \text{B2+}} \times \mathbb{1}_{t \geq \text{Entry start date}} + \Gamma X_{ct} + \epsilon_{ct} \tag{1.7}$$

in order to estimate the effect of entering eligibility to the subsidy on posterior outcomes for municipality c at time t. I also use a fully dynamic specification of the form:

$$y_{ct} = \alpha_c + \gamma_t + \sum_{\substack{k=\underline{K} \\ k \neq -1}}^{k=\overline{K}} \beta_k \mathbb{1}_{c \in \text{B2+}} \times \mathbb{1}_{t=t_0(c)+k} + \beta_{\underline{K}} \mathbb{1}_{c \in \text{B2+}} \times \mathbb{1}_{t<t_0(c)+\underline{K}} + \beta_{\overline{K}} \mathbb{1}_{c \in \text{B2+}} \times \mathbb{1}_{t>t_0(c)+\overline{K}} + \Gamma X_{ct} + \epsilon_{ct} \tag{1.8}$$

where $\mathbb{1}_{t=t_0(c)+k}$ equals 1 for treated municipalities if calendar time period t is the k^{th} period before or after $t_0(c)$, the eligibility start date for town c (the start of the *Pinel* scheme), and 0 otherwise.[58] This specification allows me to visualize any differential trends prior to gaining eligibility, giving credence to the untestable assumption of parallel post-entry trends. The underlying identification assumption is that, in the continued *absence* of an agreement, B2+ cities would have followed comparable trends to *never-eligible* B2- towns. The government's decision to classify B2 towns in a common zone suggests they initially had comparable observable dynamics. Time-varying controls, sub-sample analyses, and robustness checks support the credibility of the identification assumption.

Potential threats to identification

Selection into treatment Specific to the *entry* approach, one may be concerned about selection of B2+ towns relative to B2- cities, if locations expecting stronger investor demand were more likely to request and/or obtain an agreement. To emphasize the causal role of tax subsidies in triggering additional investment at the local level, I evidence the absence of differential pre-trends between eligible B2+ and untreated B2- areas, prior to eligibility changes. I also provide within-municipality-month placebo tests, by comparing the dynamic impact of eligibility on sales of new homes, to its placebo and null effect on sales of (ineligible) existing units. Sharp trend breaks at the beginning and end of eligibility, and high-frequency time bunching in anticipation of policy changes, also do not fit alternative explanations based on lower-frequency differential trends, and suggest a causal effect of the policy. Finally, the *exit-DiD* approach, in which the loss of eligibility is exogenous and shown to trigger immediate trend breaks in local outcomes, is an additional check allowing me to rule out the selection explanation for differential outcomes in B2+ towns during the earlier period of policy eligibility.

Other place-based government policies The validity of causal estimates obtained from the specifications above may be threatened by other time-varying determinants of municipal level-outcomes (included in ϵ_{ct}) evolving jointly with the spatial allocation and the time path of eligibility, chief among which is the presence of other place-based government poli-

[58]The continued ineligibility of non-agreed "B2-" municipalities creates a natural permanent and large control group of never-treated cities, a feature particularly relevant to the identification of the dynamic path of treatment effects (Borusyak, Jaravel, and Spiess, 2021).

cies. The only case of such a potential confounder is France's interest-free loan policy for first-time homebuyers of owner-occupied units (*Pret a Taux Zero* or PTZ). The generosity of the interest-free loan policy varied with zoning, and was amended specifically for B2 and C areas after 2017.[59] Two aspects of my empirical strategy help assuage this concern. First, using the differential treatment of B2+ and B2- municipalities upon *entry* provides compelling evidence that any effects are driven by the *Pinel* policy, since PTZ eligibility and amount do not vary between these two subsets of B2 towns. Second, I show that the impact of the *Pinel* policy is concentrated on buy-to-let units, which are ineligible to the PTZ, and high-frequency effects are only consistent with *Pinel* variation.

1.4.3 Summary of the data

I summarize the key relevant data sources used to examine the impact of the *Pinel* scheme on build-to-let sales, prices, the housing stock, local communities, landlords, and renters.[60]

Housing markets I avail myself of exhaustive data recording all housing and land transactions in France from January 2010 to December 2019, recorded by the French Treasury for tax purposes in the *DV3F* database, and comprised of 800,000-1,000,000 transactions each year over the period. I also use an exhaustive survey of developer-led projects, the *Enquete sur la Commercialisation des Logements Neufs* (ECLN), recording detailed information on all new housing sales by developers at the dwelling level.

Housing stock I use the *Sit@del2* database from the French Housing authority, comprising the universe of housing permits requested from 2009 to 2019. To study changes in the make-up of local dwellings, I rely on register data - the *Fichiers Fonciers*, an annual cadaster of all occupied and vacant housing units in France. I also use the *Répertoire des logements locatifs des bailleurs sociaux*, a repository of all social housing units in France, to study the extent of crowding-in or crowding-out of social housing by intermediate-rent housing construction

[59] After January 1st, 2018, the PTZ, a policy providing households with interest-free mortgages when they buy their first owner-occupied home, was restricted to 20 percent of the cost of first-time home purchases in the B2 and C area, down from 40 percent.

[60] Beyond the main restricted-access sources mentioned here, some surrogate sources exploited in the stylized facts characterized by section 1.2, and already mentioned then, are not described in detail here, but more information is provided on the construction of all sources in appendix 1.C.

incentivized through the *Pinel* scheme. The *Observatoire de l'artificialisation des sols* database, which combines satellite information with administrative records at the parcel level, allows me to compute the amount of land converted each year from "natural" use to residential or industrial use, in each municipality.

Investor behavior and ownership structure I make use of the FILOCOM/FIDELI database, a repository of close to thirty million housing units (as of 2019) serving as primary residences in France, to measure the evolution of the ownership network. I also use the universe of French individual tax returns in 2016, an exhaustive sample of all tax units each year (collectively known as *Fichiers POTE - panelisables*), to measure landlord status and compute rents and purchase prices for subsidized units.

Local communities Data on municipal population, as well as municipality-to-municipality annual migration flows, assembled by the French statistical institute INSEE through the Census, are used to estimate population mobility. To look at specific mobility effects among lower-income households who rent, I rely on data from the *Caisse d'Allocations Familiales*, France's main outlet for social benefit payments, recording the annual number of beneficiaries of housing allowances, as well as the number of low-income households in each municipality. When assessing whether the policy shapes local economic activity, I exploit establishment-level information from the *Repertoire des Entreprises et des Etablissements* on the number of employees and exact location of each establishment in France, aggregated at the municipality-year level.

1.5 Housing policy with home-biased landlords: results

This section analyzes the direct and indirect consequences of the place-based *Pinel* rental investment subsidy, and shows that its effects on landlords, local real estate markets, and residents, are consistent with the conceptual mechanism of imperfect capital mobility across space advanced earlier.

1.5.1 The impact of place-based subsidies on housing investment

The impact of losing eligibility to the subsidy I first examine the short-run implications of losing eligibility to the *Pinel* subsidy on purchases of eligible new dwellings from developers by individual investors, exploiting the quasi-experiment induced by its discontinuation in B2+ cities after 2018. Validating the causal policy channel and the first part of proposition 1.3.2, figure 1-9 displays the key result. It plots the quarterly coefficients β_k on a time-varying dummy for being part of the B2+ group, in the full leads-and-lags specification of equation (1.6). The dependent variable is the (log) number of buy-to-let units sold in a municipality each quarter, and the sample is restricted to B1 (always-treated) and B2+ (losing eligibility) towns. Quarterly sales follow precisely parallel dynamics throughout the duration of common eligibility to the policy, with no evidence of differential pre-trends. When the policy is discontinued, after an initial spike in the last quarter in the anticipation of its end, the number of sales drops by about twenty percent in B2+ municipalities relative to B1 towns, an economically and statistically significant impact with no apparent recovery up to eight quarters later.[61]

I assess alternative specifications and the robustness of the result in figure 1.A.2 and table 1.B.5. First, to show that the result is not driven by functional form choices or regression adjustments, panel (a) of figure 1.A.2 plots the (normalized) quarterly number of new buy-to-let units sold by developers in two specific subsets of comparable municipalities, all classified as B2 until the September 2014 re-zoning: those later upgraded to B1 by the rezoning ("B2 to B1", eligible to the *Pinel* scheme throughout the 2014-2020 period) and those that remained in the B2 group with an agreement ("B2 to B2+" towns, eligible only until December 2018). Sales of new units rose in tandem in both groups after the 2014 introduction of the *Pinel* investment tax credit, reflecting their common eligibility. However, while they remained at an elevated level in the "B2 to B1" group of "always-treated" towns, they dropped sharply and immediately once eligibility was discontinued after 2018 in "B2 to B2+" towns.

Panel (b) of figure 1.A.2 plots the quarterly coefficients β_k, where the dependent variable is the total number of units sold in a municipality each quarter and the sample is again all B2+ and B1 towns. Sales of overall new units drop after the end of the policy, although by

[61] All specifications cluster standard errors at the level of individual municipalities, to reflect potential serial correlation in error terms and the level of treatment assignment.

Figure 1-9: **The impact of exiting the subsidy on new home sales**

Figure 1-9 documents the estimates of coefficient β_k in equation 1.6 for the impact of losing eligibility to the subsidy (in B2+ cities) on (log) sales of new buy-to-let units by developers, relative to B1 (always-treated) cities. Data are computed from the ECLN database, recording exhaustive dwelling-level information on all developer sales. Standard errors are clustered at the municipality level, and 95% confidence intervals are shown in brackets. The dashed vertical line marks the end of eligibility of B2+ towns to the *Pinel* incentive relative to B1 towns.

a more limited magnitude than the specific effect on buy-to-let eligible dwellings, as predictable if some units are sold to owner-occupiers or institutional investors. Panel (c) in figure 1.A.2 shows the effect on the (log) value of sales of eligible dwellings. Finally, panel (d) examines the extensive margin response of housing investment at the town level. It plots quarterly coefficients on the B2+ dummy in a specification where the outcome is the presence of *any* new home sale in the municipality in a quarter. There is an 8-10 percent drop in the probability of a sale occurring in formerly treated B2+ towns after they lose the tax-favored treatment of new rental dwellings.

All specifications (summarized in table 1.B.5) evidence closely parallel trends (with pre-exit coefficients statistically indistinguishable from 0) for the B1 and B2+ groups throughout the period of joint eligibility, from October 2014 to December 2018; and a sharp trend break after B2+ towns lost eligibility to the *Pinel* subsidy. Overall, the results highlight the positive but quantitatively limited user cost elasticity of residential investment by individual landlords. The summary computation in table 1.B.3 estimates the net present value of the tax credit, once accounting for rent discounts and depending on assumptions on the discount rate and baseline price-to-rent ratios, ranges from 9 to 15 percent of the purchase price of a unit. Therefore, the estimated reduced-form effect is consistent with a user cost elasticity in the range of 1.3 to 2.1, markedly lower than firm-level estimates in Zwick and Mahon (2017).[62]

The impact of gaining eligibility I next turn to the *entry* approach, comparing B2+ towns (that benefited from the *Pinel* scheme from 2014 to 2018) to B2- towns, which, although part of the same class, were never granted eligibility to the incentive. Figure 1-10 plots the estimated time-varying causal effect of being part of the B2+ group on eligible new home sales to investors (equation 1.8), in the restricted sample of (eligible) B2+ and (untreated) B2- towns over the 2013-2020 period. After the implementation of the *Pinel* scheme, sales increase in eligible towns, and display a clear divergence from B2- control towns. The intensive margin impact of 20 to 25 percent at the peak is in line with the estimated effects of *exiting* the policy plotted in figure 1-9.

[62]Appendix 1.D.1 provides a discussion of the assumptions and sensitivity analyses for the implied net present value of the *Pinel* incentive, for a range of parametrizations of the discount rate, and bite of affordable rent markdowns.

Figure 1-10: **The impact of gaining eligibility to the subsidy on new home sales**

Figure 1-10 documents the estimates of coefficient β_k in equation 1.8 for the impact of gaining eligibility to the subsidy (in B2+ cities) on (log) sales of new buy-to-let units by developers, relative to B2- (never-treated) cities. Data are computed from the ECLN database, recording exhaustive dwelling-level information on all developer sales. Standard errors are clustered at the municipality level, and 95% confidence intervals are shown in brackets. The two dashed vertical lines mark (respectively) the beginning and end of differential eligibility of B2+ towns to the *Pinel* incentive relative to B2- towns.

I present alternative specifications in figure 1.A.3: the total value of sales in the B2+ and B2-; and the coefficients from estimating equation 1.8 for (log) sales of all units, and the intensive and extensive margins responses of any new sale to rental investors – respectively in panels (a), (b), (c), and (d). Panel (a) provides aggregate evidence that the effect of the policy upon entry was noticeable: sales dynamics in B2+ towns exhibit a sharp increase after the start of the scheme, while sales in B2- cities remained mostly flat throughout the duration of the policy. In each of the other three panels, the time-varying quarterly coefficient β_k on being part of the B2+ group is statistically indistinguishable from zero both before and after the differential policy eligibility period in the sample of all B2 towns, but B2+ towns exhibit a temporary departure from trend corresponding exactly to the period of availability of the investment incentive. Estimated magnitudes at both the intensive and extensive margins are summarized in table 1.B.6. They are quantitatively consistent with the estimated effect of *losing* eligibility examined *supra*.

Therefore, *entry* into eligibility to the *Pinel* policy among B2+ towns obtaining an agreement is associated with similar but opposite-sign effects to their later exit out of the policy. Beyond providing an alternative experiment to test the robustness of the findings on new dwellings sales, the similar magnitudes give reassurance that B2- cities, a never-treated permanent control group, also provide an accurate counter-factual for the trajectory that other outcomes would have followed in eligible B2+ towns, absent the subsidy.

A placebo test: existing home sales The impact of eligibility to the policy may nonetheless be confounded by differential local trends if B2+ cities exhibit systematically different housing market activity from B2- towns during the period. Using administrative records of transaction deeds from the *DV3F* database, I can compare the relative evolution of new home sales (eligible to the Pinel tax credit) and *existing* home sales, which are never subject to the landlord incentive.[63] Contrary to sales of new homes, which increase substantially in B2+ locations (figure 1.A.4, panel (a)), transactions for existing homes (panel (b)) are unaffected and follow exactly parallel trends in B2+ and B2- towns throughout the duration of the policy. This within time-city placebo test shows that the increase in developer sales in treated locations during the eligibility period is not the product of differential housing turnover, which would lead these towns to experience higher sales activity overall. On the contrary, the null effect on sales of ineligible, existing dwellings confirms the causal interpretation of the differential increase in new home sales, and supports the implicit parallel trend assumptions.

Real effects on new construction The previous results demonstrate that place-based investment tax credits substantially increase new home sales in eligible areas during the policy period. However, this behavioral response may reflect the crowding-out of un-subsidized rental housing production or owner-occupied units; tear-downs and renovations of existing dwellings; or simply faster outflows of developer inventory, rather than an effective new supply of local housing.[64] To show the limited extent of crowding-out effects, and ev-

[63]New homes are identified in the notary deeds data by their value-added tax treatment, since existing homes are not subject to the VAT.

[64]For example, Berger, Turner, and Zwick (2020) find that the First Time Homeowner Tax Credit in the United States failed to causally trigger any additional new construction, and only led to a reshuffling of the existing housing stock. Several studies of the LIHTC also find almost full crowding-out of unsubsidized con-

idence that the *Pinel* landlord incentive indeed increased the provision of housing space, I next show that the policy, rather than merely leading to faster outflows of existing developer inventory, triggered entirely new housing starts for multi-family buildings. Exploiting the universe of individual building license requests from 2010 to 2020, I examine building license requests, as well as authorized new units and housing starts in targeted B2+ areas relative to excluded B2- towns, during the policy period. Figure 1.A.5 summarizes the results.

Its top panels plot the quarterly number of authorized multi-family units in B2+ (once-treated) and B2- (control) or B1 (always-treated) municipalities, normalizing counts by the pre-period average quarterly number of authorizations. Licenses granted in B2+ and B2- municipalities closely track each other before the start of the policy, even at high-frequency seasonal variation. After the policy starts, however, the number of licenses granted takes off in eligible cities relative to the trend in ineligible cities, with a noticeable divergence throughout the year 2017, up until December 2017. Sharp spikes occur in the months surrounding the end of construction eligibility in December 2017, a natural consequence of time-bunching in licenses granted to corporate developers evidences earlier (see appendix 1.F.1 for additional evidence on such time-bunching). After eligibility ends for building license requests in B2+ towns, the trend in these formerly eligible locations rapidly falls back to the corresponding path for ineligible cities. The two series come back in lockstep as the differential subsidy to eligible areas ends, and exhibit similar trajectories in 2019-2020. This differential variation in authorized units, closely mirroring the timing of policy eligibility to the subsidy, implies that the *Pinel* subsidy was indeed responsible for the rise in authorized constructions in eligible municipalities.

Panel (c) confirms this pattern using the cumulative monthly number of *housing starts* in both types of B2 municipalities. Housing starts follow a similar qualitative pattern to housing units authorized. After following precisely parallel trends prior to the start of the Pinel scheme in September 2014, housing starts diverge shortly after the subsidy is put in place in eligible B2 cities, relative to their ineligible counterparts. Panel (d) estimates the full leads-and-lags specification for the evolution of licenses granted in B2+ cities relative to B2- locations at annual frequency, finding a close to ten percent increase in licenses at the 2017

struction by subsidized units.

peak.

Urban sprawl and land conversion New home construction may occur either through the demolition and replacement of the existing stock of housing; through increased density of new floor area built per square footage of land in already urbanized spaces; or through expansion of the urban area into formerly un-built plots of land, leading to urban sprawl. To evidence the impact of the policy on urban sprawl and the share of local land that is built-up, I avail myself of data combined by the French *Observatoire de l'artificialisation des sols*, reporting detailed measures of land use at the local level, notably the annual flow of land area converted from natural uses (forest or agricultural land) to new habitat. Figure 1.A.6 shows the annual flow of land transformed from natural uses to residential uses ("artificialization"), across municipalities, depending on their eligibility status. It evidences differential dynamics in Pinel-eligible and ineligible areas, with a pronounced divergence occurring in 2016, when conversion of land for residential use picked up in B1 and eligible B2 municipalities, but remained low in ineligible B2 and C towns. The differential spike in 2016 is consistent with the timing of the rise in construction observed in eligible areas, relative to ineligible towns. The policy therefore led not only to increased new construction of housing in targeted locations, but also to substantial urban sprawl and conversion of land from natural to residential uses at the boundary of urban areas, a potential driver of increased commuting times and adverse environmental consequences.

Aggregate effects of landlord subsidies on sales In line with the model-consistent mechanism of non-housing to housing investment reallocation, I provide suggestive evidence that aggregate rental investment by individual owners responds to time-varying tax incentives, exploiting longer-run aggregate time series data and all return shocks induced by a more or less favorable tax treatment. Even before the *Pinel* scheme, French governments have provided tax subsidies to the individual purchase of newly built rental housing, starting in the 1990s – generally conditional on setting rents at an intermediate level between market-rate and public housing. Table 1.B.4 summarizes the successive schemes in operation from 1996 to 2021, noting whether the implementation of a new scheme resulted in increased or decreased support relative to earlier tax subsidies.[65]

Figure 1-11 shows new residential units purchased each quarter closely track the implemen-

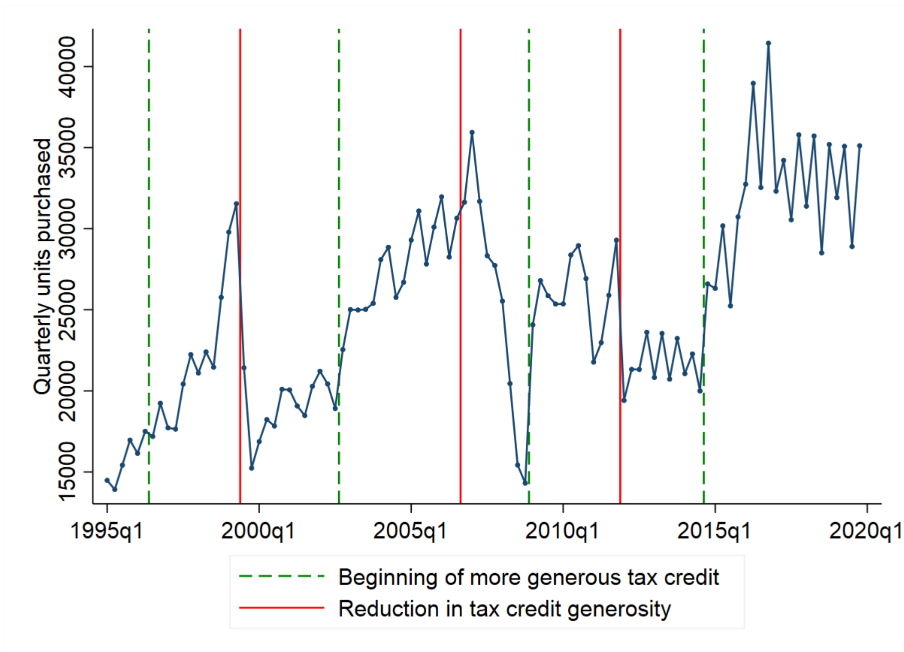

Figure 1-11: **New units sold by developers, quarterly**

The figure plots the total quarterly number of new housing units purchased from developers in France from 1995 to 2019. Dashed vertical green lines indicate the beginning of a more generous investment tax credit or bonus depreciation scheme for new housing; continuous vertical red lines indicate a switch to a less generous tax regime, either through reduced tax credit levels, stricter affordability requirements, or a deceleration of bonus depreciation schedules.The corresponding successive tax regimes are described in detail in table 1.B.4 and appendix 1.D.1. The *Pinel* scheme starts with the last green line, in 2014Q4. The total number of units sold is computed from microdata in the *ECLN* database, an exhaustive survey of developer-led projects of five or more units.

tation of more generous subsidy schemes, suggesting a substantial aggregate responsiveness of rental investment to tax incentives.[66] Periods in which incentives were altered to be more favorable to investors all coincide with a sharp rise in the number of units sold by devel-

[65] Appendix 1.D.1 provides additional descriptive information on the bonus depreciation and investment tax credits subsidizing new dwellings purchases by individual investors since 1996. Some resulted in a higher net subsidy to rental investment: the 1996-1998 initial *Perissol* accelerated depreciation scheme, the 2003-2005 mostly unconditional *Robien* tax credit, the 2009-2010 stimulus-motivated *Scellier* Act, and, finally, the 2014 flexible-duration *Pinel* scheme). Others reduced the net benefit to landlords: the 1999-2002 *Besson* law slowing down accelerated depreciation provisions, the 2006-2009 "re-focused" *Robien* scheme, the 2011-2012 budget-related cuts to the level of the *Scellier* tax credit, and the 2013 *Duflot* law with its more stringent affordability conditions.

[66] Figure 1.G.6 plots the corresponding time series for the inflation-adjusted euro *value* of new residential investment.

opers. Conversely, intervals with lower subsidies are accompanied by reduced sales.[67] The generosity of these successive tax credits reflects changing government priorities and endogenously responds to the state of the macro-economy, therefore requiring the difference-in-differences approach exploited in the core quasi-experiment of the paper to make causal claims. Nonetheless, while some of the variation, such as the 2009 trough, reflects business cycles regularities, high-frequency changes precisely track the evolution of tax credits, are oftentimes at odds with the macroeconomic cycle,[68], and strongly suggest the presence of aggregate effects of landlord subsidies on investment. The large observed rise in sales after the implementation of the *Pinel* scheme in October 2014, from c. 20,000 quarterly sales to c. 35,000, is consistent with the substantial place-specific causal effects estimated earlier not cancelling in the aggregate.

Moreover, official deeds data allow me to document anticipation effects and inter-temporal tax arbitrage in aggregate sales. Figure 1.A.7 documents the high-frequency pattern of actual signing dates for all new homes sales deeds recorded by notaries. Substantial time-bunching occurs in the weeks immediately preceding an anticipated shift to less favorable tax incentives. For example, while sales made before December 31st 2011 were eligible to the 2011 Scellier subsidy of 22 percent (for units respecting environmental guidelines – 13 percent otherwise), sales posterior to that date were only granted a lesser tax credit of 13 percent (6 percent for non-energy efficient units). Consequently, a marked bunching of sales occurs in the last days of 2011, immediately before the switch to the less generous subsidy.[69] Similarly, end-of-fiscal year spikes in recorded deeds are visible after the implementation of the *Pinel* subsidy scheme in 2014. This abnormal end-of-year concentration of transactions is likely triggered by increased certainty about a household's precise income tax liability (and therefore the ability to collect the full non-refundable *Pinel* tax credit), as in Xu and Zwick (forthcoming)'s real option model of fourth-quarter corporate investment spikes.

[67] For the shorter, post-2010 period in which direct information on the final use of the asset is available, figure 1.G.7 additionally evidences that the aggregate share of buy-to-lets in new units sales mirrors the time variation in subsidy generosity. It sharply drops at the end of the more generous version of the *Scellier* scheme in 2011, and rises after the implementation of the *Pinel* tax credit.

[68] e.g. during the dot-com boom and bust, or after the 2014 implementation of the *Pinel scheme*

[69] A tolerance period until March 2012 (for sales agreed with developers before December 2011) is also accompanied by a visible spike in deeds recorded immediately before March 31st, 2012, and a missing mass of sales in the following weeks. This anticipation result confirms the findings of Singh (2019), who documents that developers bunch residential investment decisions immediately before eligibility deadlines, in expectation of lower future property tax incentives.

1.5.2 Home bias and imperfect capital mobility: direct evidence

Incomplete price capitalization

According to proposition 1.3.2, reducing the subsidy to rental units in location j is likely to decrease the price of new housing, as demand from investors dries down. The impact of losing eligibility to the subsidy on new unit prices is shown in figure 1-12, plotting estimates from equation 1.6, taking the log price of new dwellings sold by developers as the outcome variable. Quantitatively, the effect is however limited: table 1.B.5, column (4) shows that prices drop by c. four percent after the end of eligibility. Under standard perfect capital mobility models, required returns to housing investment are given by an exogenous interest rate determined in global financial markets. As a consequence, a place-based housing subsidy to landlords in one location should be accompanied by a full capitalization of the value of the incentive in purchase prices.[70]

Computations of the value of the subsidy in table 1.B.3 (under the null of perfect capital mobility, and for a range of assumptions for the discount rate, baseline rent-price ratio, and bite of affordability requirements) imply a minimum value of 8 to 9 percent (for high discount rates and large potential discounts from rent ceilings relative to market values) and a more central value of 12 to 15 percent of the purchase price (for more standard discount rates and affordability discounts consistent with observed values). However, the observed capitalization effects in new unit prices reach a maximum of about four to five percent after several quarters, only consistent with imperfect capitalization. This provides evidence for a key prediction of proposition 1.3.2: to draw additional investment to target areas, place-based policies dS_j must raise the net return r_j^H to rental assets in these cities. This implies that part of the incidence falls upon infra-marginal savers. Specifically, because most nearby investors are infra-marginal in their investment choices, due to the presence of home bias, a place-based tax credit benefits local investors, while higher returns are required to make long-distance landlords break even – in terms of full returns, inclusive of their distance costs.

[70] See figure 1-7 and the discussion of proposition 1.3.2.

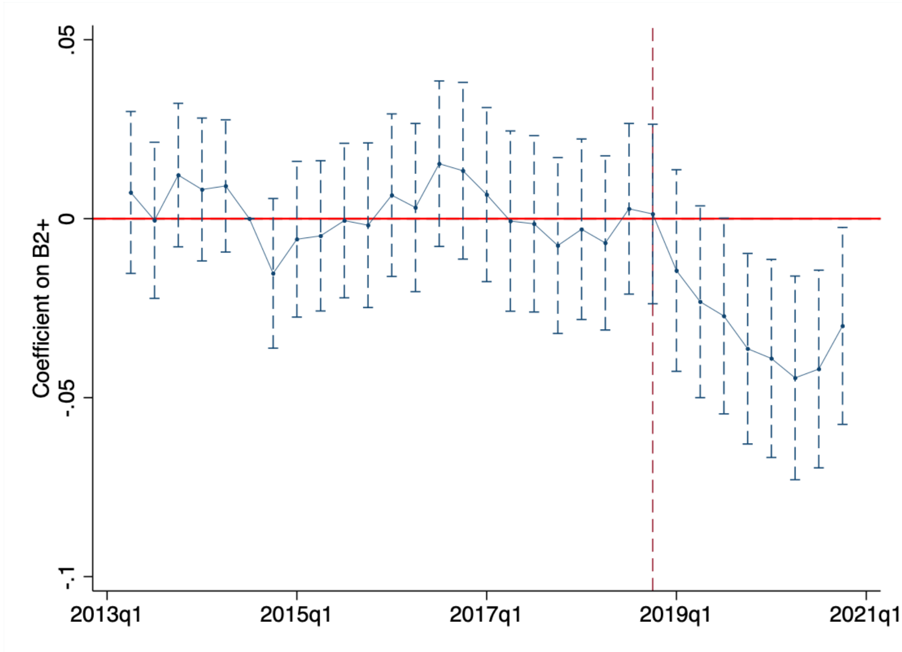

Figure 1-12: **The impact of exiting the subsidy on new home prices**

Figure 1-12 documents the estimates of coefficient β_k in equation 1.6 for the impact of losing eligibility to the subsidy (in B2+ cities) on (log) prices of new buy-to-let units by developers, relative to B1 (always-treated) cities. Data are computed from the ECLN database, recording exhaustive dwelling-level information on all developer sales. Standard errors are clustered at the municipality level, and 95% confidence intervals are shown in brackets. The dashed vertical line marks the end of eligibility of B2+ towns to the *Pinel* incentive relative to B1 towns.

Increase in long-distance investment

The spatial investment decision undertaken by landlords weighs the returns of a given residential investment against its (spatially varying) costs - inclusive of any subsidy. As subsidies offset part of the costs, one key prediction of the model is that the reallocation effects on housing investment are stronger among long-distance buyers, who expand the scope of locations in which they consider investing. In essence, long-distance buyers are more likely to be marginal in their decision to purchase rental assets in j. I test for this impact of buy-to-let investment subsidies on the reduction in home bias. I combine administrative information on the residence location of investors, and the construction year of rental units, for all housing units built over the 2010-2018 period in B2+ (treated) and B2- (control) towns.

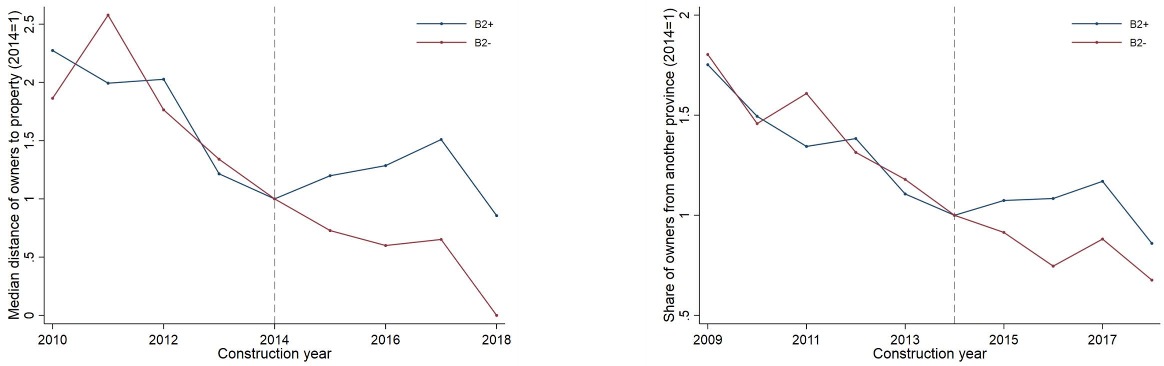

(a) Median distance

(b) Share from another province

Figure 1-13: **The spatial ownership effects of entering eligibility**

Figure 1-13 illustrates the spatial reshuffling of housing investment as a consequence of the subsidy. Panel (a) (resp. panel (b)) plots, for B2+ (treated) and B2- (control) towns, the normalized median distance of owners (resp. the normalized share of investors from another province) of rental units in municipalities in each group, depending on the construction year of the rental unit. Statistics are computed from individual ownership data from the 2019 *FIDELI* database.

Figure 1-13 displays the main result on the impact of place-based subsidies on spatial investment behavior. It compares the residence location of *landlords* for units built in B2+ and B2- locations, depending on the construction year. In particular, I test for the differential evolution, in eligible locations relative to their untreated counterparts, of the median distance of investors from their properties (panel (a)) share of units owned by landlords located in another province (panel (b). Both outcomes follow similar patterns in the two subsets of B2 municipalities for units built before the divergence in tax credit eligibility, with an increasing home bias and a downwards trend in the presence of remote landlords. However, in B2+ locations that gain eligibility to the *Pinel* tax credit after 2014, the trend reverses. The increased presence of investors coming from different provinces or living far away in these treated locations is consistent with the model prediction (part (ii) of proposition 1.3.2). Long-distance investors reallocate their housing portfolio towards cities that represent a low initial share of investments from their origin location, leading on average to a decrease in landlord home bias in locations eligible to the subsidy. This reshuffling of the spatial ownership network is a key mechanism at play in the spatially heterogeneous response of residential investment. By drawing landlords from further away, subsidies partly offset home bias, at the expense of an implicit windfall gain for infra-marginal owners located nearby who would have chosen to invest locally, even absent the subsidy.

Heterogeneous effects on the housing stock

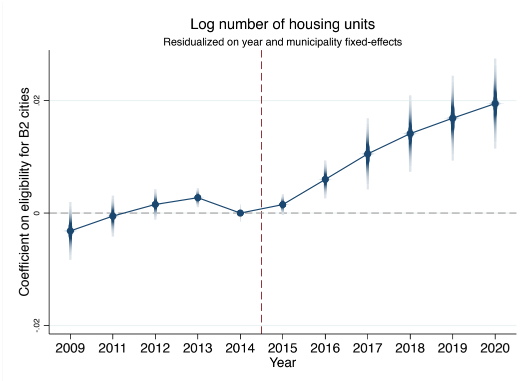

(a) Total stock

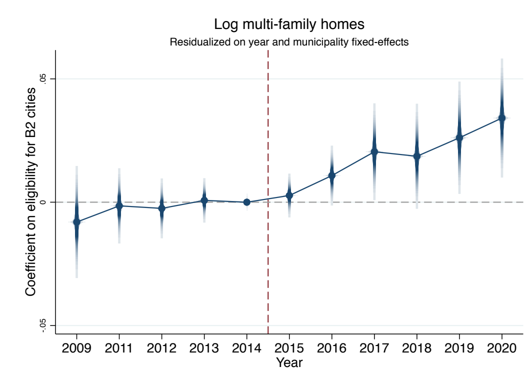

(b) Multi-family units

Figure 1-14: **The impact of the subsidy on the housing stock**

Figure 1-14, panel (a) (resp. panel (b)) documents the estimates of coefficient β_k in equation 1.8 at annual frequency, for the impact of gaining eligibility to the subsidy (in B2+ cities) on the log number of total (resp. multi-family) housing units in the municipality, relative to B2- (never-treated) cities. Data are computed from the *Fichiers Fonciers* exhaustive cadaster. Standard errors are clustered at the municipality level and 95% confidence intervals are shown. The dashed vertical line marks the beginning of eligibility of B2+ towns to the *Pinel* incentive relative to B2- towns.

Quantity and composition effects on the local housing stock: baseline effects Using unit-level annual data on all dwellings in each French municipality, I study the impact of the *Pinel* landlord subsidy on the local housing stock in subsidized B2+ locations. Figure 1-14, panel (a) shows the baseline result, in a fully dynamic difference-in-differences specification: the total stock of housing units in eligible areas increased substantially, by approximately 2 percent in the long-run, after towns gained eligibility to the policy, therefore implying no significant crowding out of un-subsidized units. The make-up of the local housing stock is substantially transformed, with a rise in the share of denser multi-family units, which corresponds more frequently to denser, smaller apartments. Panel (b) shows that the impact of the subsidy is concentrated on multi-family units, with a rise of close to 4 percent in the medium-run; on the contrary, the policy has no effect on the growth of single-family units in treated municipalities. Because the housing market tends to be highly segmented between single-family, mostly owner-occupied housing, and multi-family, generally renter-occupied units, the subsidy to rental investment generates a shift in the composition of the local housing stock, with the apartment share increasing sharply.[71]

Heterogeneous effects by openness to outside capital The third part of proposition 1.3.2 predicts that, as a consequence of the stronger response of distant investors, the effect of the subsidy on the local housing stock will be stronger in locations more open to outside capital. I estimate the heterogeneous effects of the subsidy on the housing stock, using all B2 locations, according to the following specification:

$$y_{ct} = \alpha_c + \gamma_t + \beta_R \mathbb{1}_{c\in \text{B2+}} \times \mathbb{1}_{t\geq \text{Entry start date}} + \theta_R \mathbb{1}_{c\in \text{B2+}} \times \mathbb{1}_{t\geq \text{Entry start date}} \times S_c + \Gamma X_{ct} + \epsilon_{ct} \tag{1.9}$$

where S_c denotes the dimension of heterogeneity at the level of municipality c. I define more "capital-open", locations c, consistent with the model-implied measures described in proposition 1.3.2, as those situated in a commuting zone i exhibiting a lower concentration of ownership $\sum_i \frac{H_{i\to j}}{H_j} s_{ij}$, where s_{ij} is the share of housing assets in j in the housing portfolio of investors from i. Since the model also implies that a higher land share in housing production by developers lowers the effective physical supply elasticity of housing, locations with a production technology more reliant on the fixed local supply of land exhibit lower responses. While measuring parameters of the local production technology is challenging, I control for the heterogeneous response of cities by interacting the post-policy indicator with the extent of local buildable land constraints, using the share of urbanized land in the municipality (gathered from satellite information) as a proxy for the bite of land constraints in the spirit of Saiz (2010).

The results, presented in Table 1.B.7, provide evidence that the policy witnesses stronger effects in towns that are more open to external capital, as characterized by a smaller CZ-level ownership concentration term $\sum_i \frac{H_{i\to j}}{H_j} s_{ij}$. Across specifications (which include time fixed effects interacted with baseline municipal characteristics, or weighting by population), while the baseline effect of the policy is a substantial medium-run increase in the housing stock of c. two percent, a higher concentration of ownership in the hands of nearby investors

[71] Figure 1.F.6 in appendix 1.F.2 evidences an additional induced effect on the local number of social housing units.

limits the treatment effect. Relying more on local investors indicates that a municipality is more closed to outside capital, and therefore less exposed to the potential reallocation effects of the policy on the portfolio of long-distance owners, a central mechanism through which subsidies raise the local housing stock.

1.5.3 The impact of capital reallocation on worker mobility

Consistent with the theoretical predictions, landlord subsidies increase sales, construction activity, and the local housing stock in targeted areas. A larger local housing stock, however, may lead mostly to an increased per capita consumption of housing, if most of the effect operates through local family formation and household splitting. On the contrary, a last key prediction of the model lies in the fact that the higher local supply of private rental housing induced by tax credits and subsidies affects population growth and inwards mobility.

Using town-level data on population and mobility decisions, I examine the impact of the additional private housing induced by the tax subsidy on inwards population movements and overall town population. Panel (a) of figure 1-15 graphically presents the results of the estimation specification (1.8), where the outcome variable is the total population of residents in a town. It plots the coefficients of interest β_k (at annual frequency) corresponding to the differential time-varying impact of being part of the B2+ treated group, relative to untreated B2- cities. In order to focus on medium-term dynamics and because of the low frequency of Census re-sampling, I restrict treated cities to those B2+ towns that had obtained an agreement before the start of the *Pinel* scheme in 2014. The empirical results show that the scheme leads to a slightly more than 1 percent increase in population in treated towns over the long-run, a slightly smaller effect than the housing stock impact.

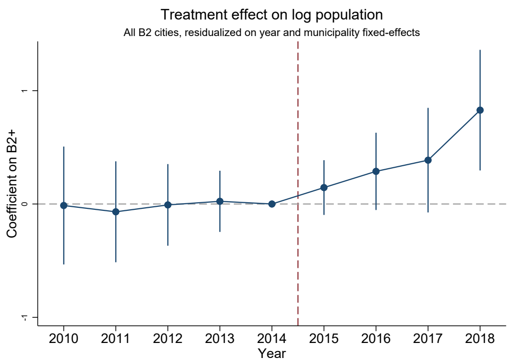

(a) Population

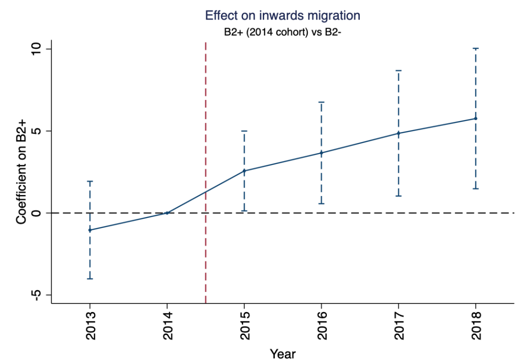

(b) Inwards migration

Figure 1-15: **The spatial reshuffling of people**

Figure 1-15 illustrates the spatial reshuffling of population as a consequence of the landlord subsidy. Panel (a) (resp. b) plots, for B2+ (treated) relative B2- (control) towns, the coefficient on eligibility in equation 1.8 for the log population (log in-migrants number) in the town. Standard errors are clustered at the municipality level and 95% confidence intervals are shown in brackets. Statistics are computed from individual Census data.

The composition of the local population also evolves towards lower-income tenants as a consequence of the policy. Using administrative data from the *Caisse d'Allocations Familiales* on social benefits recipients at the town-level, I find an increase in the number of recipients of housing allowances in targeted B2+ locations after the implementation of the policy. Figure 1.A.8 evidences the differential evolution of the number of recipients of housing vouchers in B2+ and B2- cities over time after 2014, and shows that, while the number of owner-occupier recipients of benefits did not increase, the number of tenants eligible to housing benefits displays a clear rise in B2+ locations relative to B2- controls after the start of the policy in 2014. Finally, to further investigate the role that spatial mobility plays in triggering additional migration towards treated locations, I show that differential inwards migration trajectories between targeted and untargeted municipalities account for the effect on population. Panel (b) of figure 1-15 graphically presents the results of the estimation specification (1.8), taking as the outcome variable the log number of people who moved into a locality within the last year. While data on the origin of residents in the past year is only available starting in 2013, the event-study results demonstrate that, after following parallel trends prior to the policy introduction, B2+ towns treated after 2014 by the landlord subsidy witness a sharp increase in annual flows of inwards migration, drawing additional households and accounting for most of the medium-term rise in population. Consistent with the model's implied mechanism, shifting the allocation of landlords' portfolios through tax policies induces increased in-migration effects on tenants, and a rise in the population of targeted locations.

1.6 Conclusion

This paper shows that frictions constraining landlord investment behavior across places lead to home-biased landlords, and matter for housing affordability. In a spatial world of footloose renters and immobile landlords, the location of owners themselves shapes the affordability of rental housing. The physical proximity of landlords to a location acts as a force to lower "financial trade costs" in the provision of residential capital and rental services. This agglomeration force operates not through the production (Marshallian knowledge effects) or consumption (scale effects) of the final good, but through financial frictions for long-term capital in a spatially segmented housing market.

Precisely because of spatial barriers specific to residential investment, the geography of *investors* matters for the response of their portfolio allocation to place-based tax incentives, and for the ability of housing supply to accommodate increased demand to live in an area. I estimate empirically the response of residential investment to return shocks, using quasi-experimental variation in a tax credit for affordable new rental housing in targeted French municipalities. The policy not only hastens the sale of existing developer inventory, but also triggers new dwellings construction and purchases, with only limited crowding-out effects on owner-occupiers. The scheme increases mobility towards targeted locations, while shifting the makeup of the local housing stock, and raising the income diversity of eligible communities. Speaking directly to a partial offset of landlord home bias, the subsidy reshuffles the spatial asset ownership network across cities by drawing more remote landlords into the local rental market, and by raising the net return to investment in treated locations.

As housing costs rose steeply in dense urban agglomerations over the last three decades, governments have responded to this trend by intervening in real estate markets, and, in particular, by using tax policy to encourage landlords to provide additional housing in expensive locations. Overall, this paper demonstrates theoretically and empirically that such policies targeting affluent individual landlords can affect the quantity and the allocation of new housing supply. Governments, however, incur large budget costs in the process of implementing these place-based subsidies, because a substantial share of the incidence of these benefits accrues to infra-marginal investors, in particular those whose residence is closest to targeted locations. Whether equity-efficiency trade-offs nonetheless justify locally targeted support to the provision of housing by individual landlords, especially in the presence of heterogeneity in productivity and amenity agglomeration benefits, is a potentially fruitful avenue for future research.

Appendices

1.A Additional main figures for Chapter 1

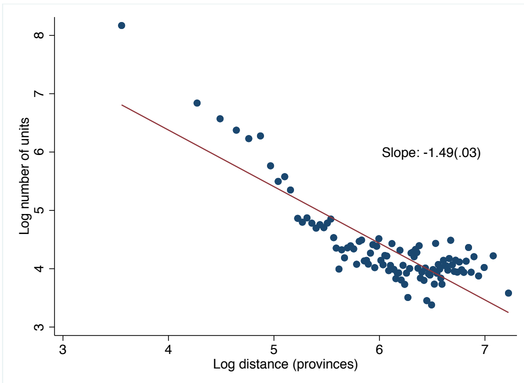

(a) Across provinces

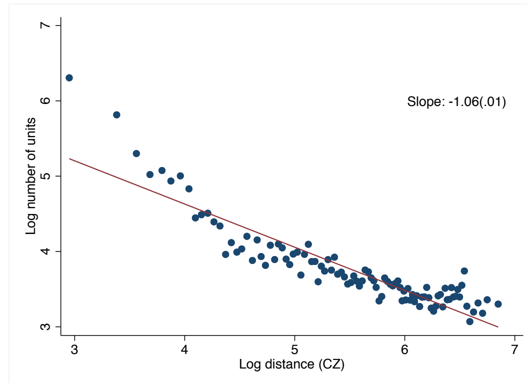

(b) Across commuting zones

Figure 1.A.1: **The home bias in homes: additional evidence**

Figure 1.A.1 documents additional evidence on the home bias in rental property investment across various aggregation levels. Panel (a) (resp. (a)) plots a binned scatter plot of the log number of units owned in province (resp. CZ) j by investor from province (resp. CZ) i against the bilateral distance between the two provinces (resp. CZ). The distance across provinces is the road distance, while the distance across CZ is the haversine distances between their main municipalities. The own-distance is computed as the average distance within a circle with radius equal to the average across CZs or provinces. See table 1.B.1 for a description of the data and underlying regressions.

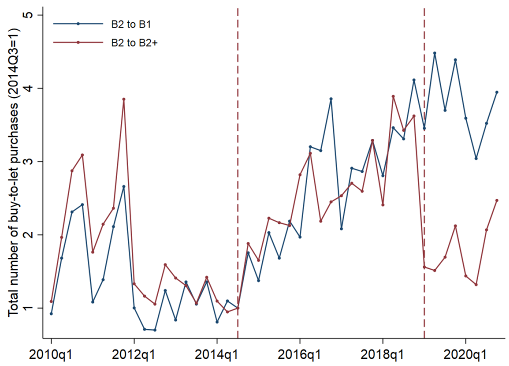

(a) Raw series

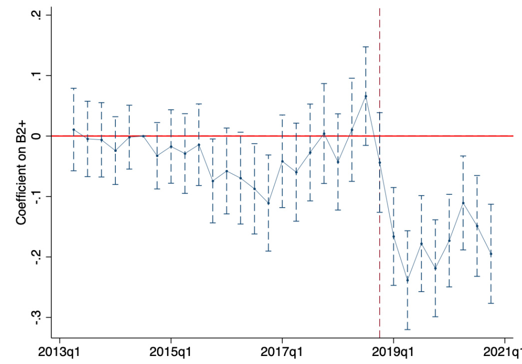

(b) Intensive margin: all new home sales

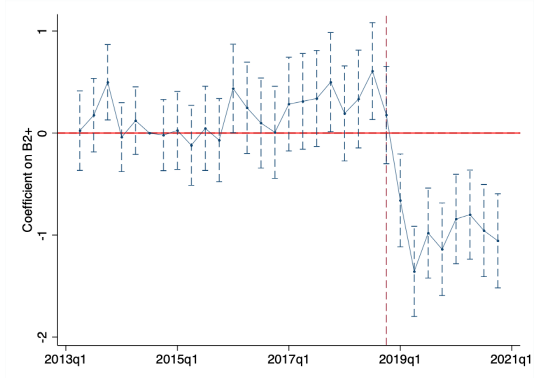

(c) Intensive margin: buy-to-let, log value

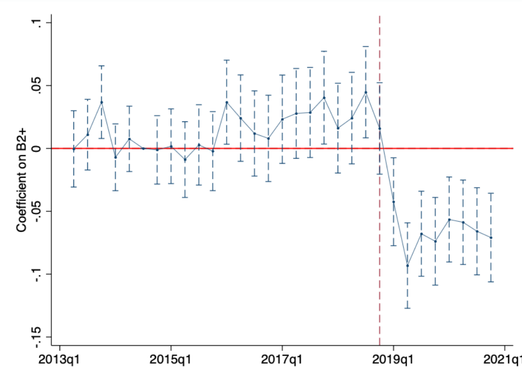

(d) Extensive margin: buy-to-let

Figure 1.A.2: **The impact of exiting the subsidy on new home sales: additional evidence**

Figure 1.A.2 documents additional evidence on the impact of losing eligibility to the subsidy (in B2+ cities) on new home sales by developers, relative to B1 (always-treated) cities. Panel (a) plots the normalized raw series of quarterly new home sales in B2+ and B1 cities among municipalities classified as B2 prior to the September 2014 overhaul. The first dashed vertical line indicate the start of the Pinel scheme; the second dashed line marks the end of eligibility for *sales* in B2+ towns. Panels (b), (c), and (d) respectively plot estimates of equation 1.6 for (log) total new home sales, the (log) value of buy-to-let sales, and the probability of any home sale. Data are computed from the ECLN database, recording exhaustive dwelling-level information on all developer sales. Standard errors are clustered at the municipality level. The dashed vertical line marks the end of eligibility of B2+ towns to the *Pinel* incentive relative to B1 towns.

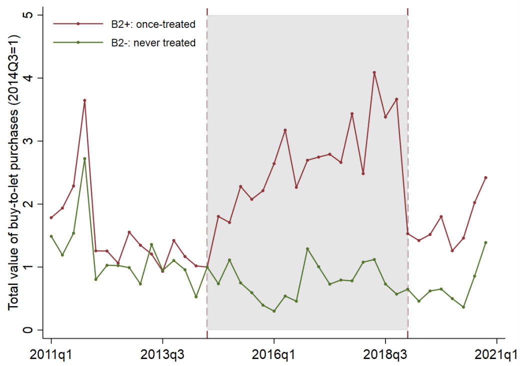

(a) Raw series

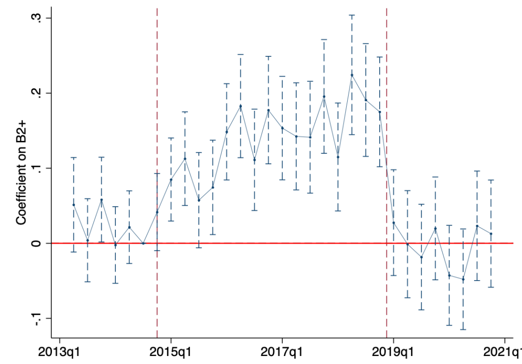

(b) Intensive margin: all new home sales

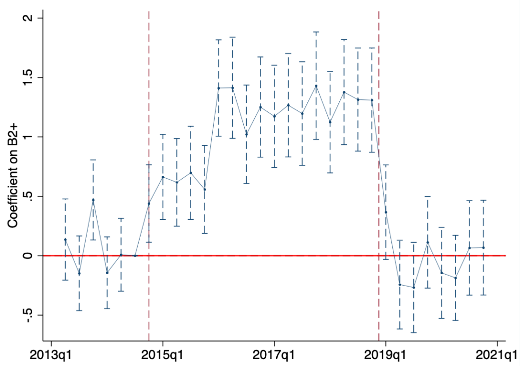

(c) Intensive margin: buy-to-let, log value

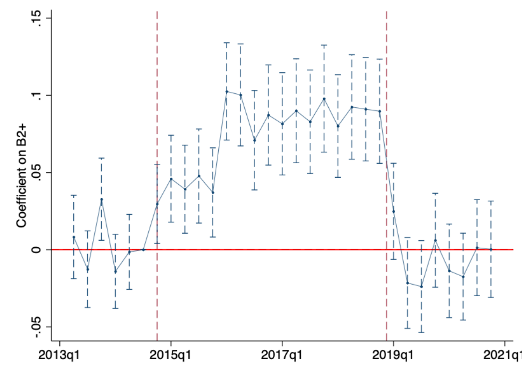

(d) Extensive margin: buy-to-let

Figure 1.A.3: **The impact of gaining eligibility to the subsidy: additional evidence**

Figure 1.A.3 documents additional evidence on the impact of gaining eligibility to the subsidy (in B2+ cities) on sales of new homes by developers, relative to B2- (never-treated) cities. Panel (a) plots the normalized raw series of quarterly new home sales in B2+ and B2- cities among municipalities classified as B2 prior to the September 2014 overhaul. Panels (b), (c), and (d) respectively plot estimates of equation 1.8 for (log) total new home sales, the (log) value of buy-to-let sales, and the probability of any home sale. Standard errors are clustered at the municipality level. The first dashed vertical line indicate the start of the Pinel scheme; the second dashed line marks the end of eligibility for *sales* in B2+ towns. Data are computed from the ECLN database recording individual information on all developer sales.

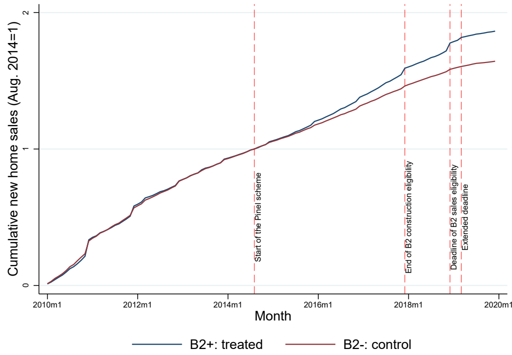

(a) Sales of new homes

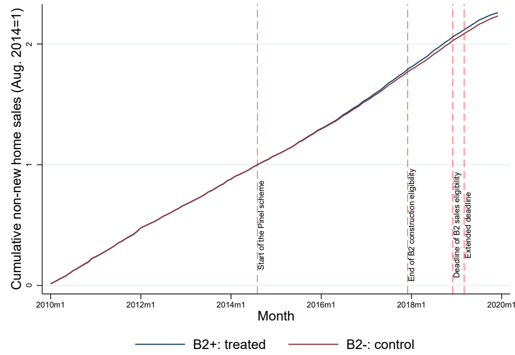

(b) Sales of existing homes

Figure 1.A.4: **Placebo test: new and non-new home sales**

The figure plots the cumulative value of home sales (normalized prior to the start of the *Pinel* scheme), respectively in never-eligible B2- and once-eligible B2+ municipalities. Panel a plots the case of new home sales, which were eligible to the landlord incentive, while panel b plots the value of non-new home sales, which were not. Dashed vertical lines mark the beginning, initial planned end, and actual end of the scheme in B2+ towns. Source: *DVF* database, an exhaustive record of housing transactions in France made available by the French Finance Ministry. Sales are restricted to houses and flats with an available transaction value.

(a) License requests, B2+ v. B1

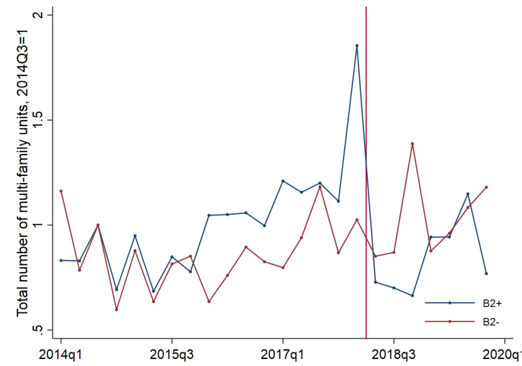

(b) License requests, B2+ v. B2-

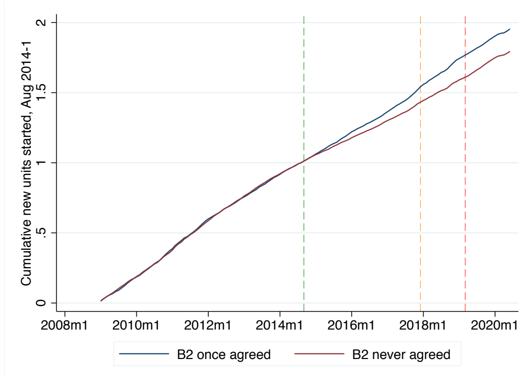

(c) Cumulative housing starts, B2+ v. B2-

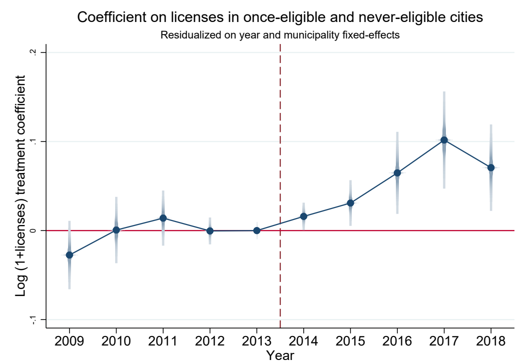

(d) Licenses granted: event-study

Figure 1.A.5: **The impact of gaining eligibility on multi-family housing licenses**

Figure 1.A.5 documents additional evidence on the impact of gaining eligibility to the subsidy (in B2+ cities) on the production of new housing, relative to B2- (never-treated) cities. Panels (a) (resp. (b)) plot the quarterly number of licenses *requested* for multi-family units in B2+ cities against B1 (resp. B2-) towns. Panel (c) plots the cumulative monthly number of *housing starts*, respectively in never-eligible B2- (red line) and once-eligible B2+ (blue line) municipalities. The green dashed vertical line indicate the start of eligibility for constructions in B2 municipalities; the orange dashed vertical line marks the end of eligibility for *licenses* , while the dashed red line marks the end of eligibility for *sales*. Panel (d) plots yearly coefficients on eligibility to the *Pinel* incentive for the log number of licenses *granted*, estimating the differential departure from trend in B2+ (relative to B2-) cities. Standard errors are clustered at the municipality level. Sources: *Sit@del2* database, an exhaustive municipality-level repository of all licenses granted - see appendix 1.C for a detailed description of the data.

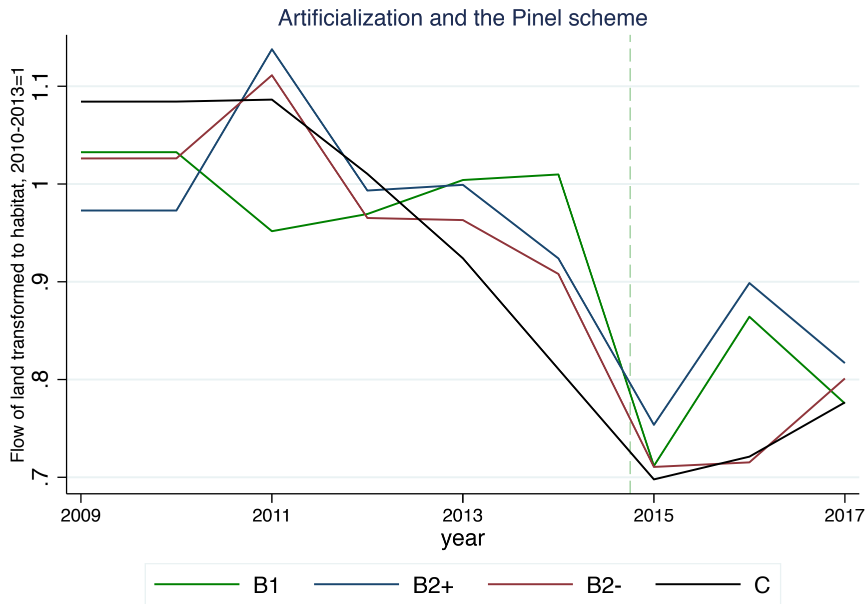

Figure 1.A.6: **Total land area shifted from natural to residential uses**

The figure plots the annual area of land converted from natural to residential use in France from 2009 to 2017, respectively in B1 (green line), once-eligible B2 (blue line), never-eligible B2 (red line), and never-eligible C (green line) municipalities. The annual flow of land converted is normalized to the average of the period immediately preceding the implementation of the policy (2009-2014). Sources: *Observatoire de l'artificialisation des sols* database - see appendix 1.C for a detailed description of the data.

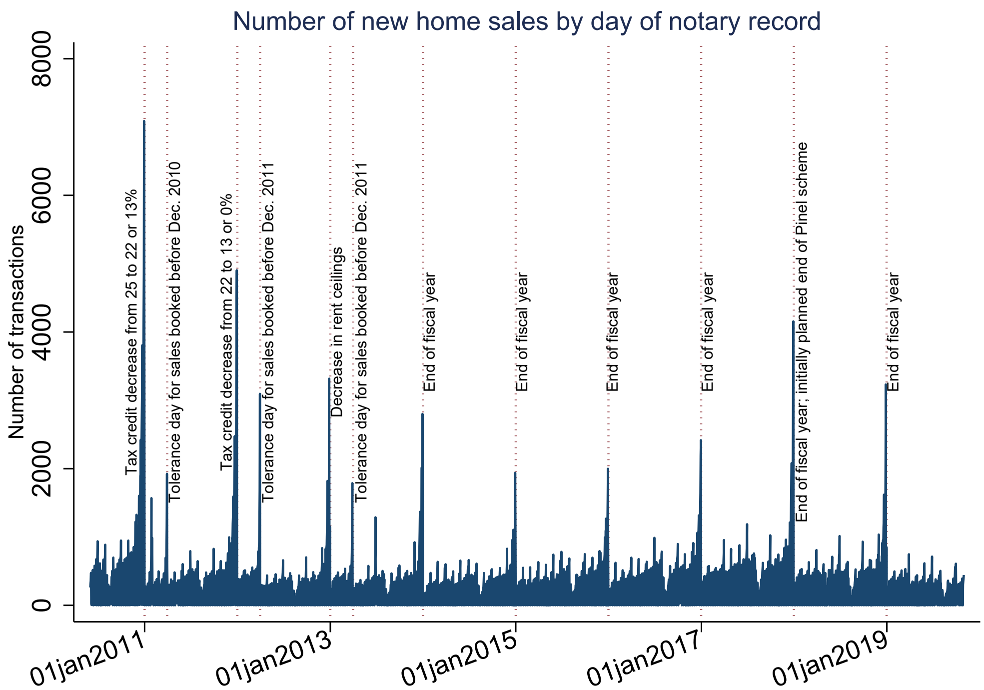

Figure 1.A.7: **New home sales recorded by notaries, weekly**

The figure shows the total weekly number of new housing units sold in France from 2010 to 2019. Dashed lines indicate eligibility deadlines (or deadline extensions) before the switch to a less generous tax regime, as well as the end of each fiscal year. The successive tax regimes are described in table 1.B.4 and appendix 1.D.1. The total number of units sold is computed from microdata in the *DV3F* database, an exhaustive registry of all housing deeds in the country. The "new" unit status is inferred from the VAT treatment of the dwelling.

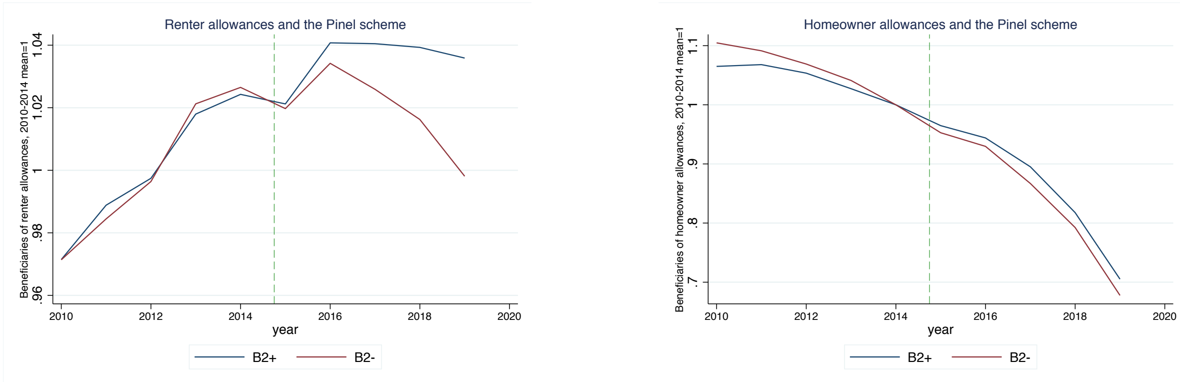

(a) Renters

(b) Owner-occupiers

Figure 1.A.8: **The evolution of recipients of social benefits**

Figure 1.A.8 illustrates the evolution of the local population as a consequence of the landlord subsidy. Panel (a) (resp. b) plots, for B2+ (treated) relative B2- (control) towns, the normalized number of recipients of social benefit allowances for rental housing occupiers (resp. owner-occupiers) in the town. Statistics are computed from *Caisses d'Allocations Familiales* municipality-level data.

1.B Additional main tables for Chapter 1

Table 1.B.1: The home bias in homes: gravity estimates

	Across provinces				
Distance coefficient	-1.206****	-1.489****	-1.241****	-0.139****	-0.108****
	(0.0277)	(0.0120)	(0.0109)	(0.0116)	(0.0109)
Twoway origin-destination FE	No	Yes	Yes	Yes	Yes
Includes *ii* pair	Yes	Yes	No	No	Yes
SCI controls	No	No	No	Yes	Yes
R-Square	0.303	0.815	0.843	0.944	0.954
Observations	7722	7721	7625	7625	7721
	Across commuting zones				
Distance coefficient	-0.573****	-1.063****	-0.779****	-1.158****	
	(0.0107)	(0.0111)	(0.0519)	(0.00709)	
Twoway origin-destination FE	No	Yes	Yes	Yes	
Includes *ii* pair	Yes	Yes	Yes	No	
Method	OLS	TWFE	PPML	TWFE	
R-Square	0.138	0.440	0.2446	0.699	
Observations	17976	17959	17959	17655	

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Table 1.B.1 estimates the impact of bilateral distance on the log number of units $\mathcal{C}_{ij}$ owned by investors from location i in location j. It reports the coefficient β on (log) distance in two-way fixed effects regressions of the form $\log C_{ij} = \alpha_i + \gamma_j + \beta \log d_{ij} + \eta X_{ij} + \epsilon_{ij}$ where X_{ij} is a (potentially empty) control for the social connectedness index (Bailey et al., 2018) between two locations. The top panel plots the estimated effect across provinces, while the bottom panel estimates the effect of log distance on log ownership across commuting zones in mainland France. To comply with disclosure confidentiality requirements, I aggregate rental dwelling ownership at either the commuting zone level (across CZ or "*zones d'emploi*") in mainland France or at the province level (across provinces or "*departements*" in mainland France), and exclude bilateral pairs of origin-destination locations with ten units or fewer.

Table 1.B.2: Summary statistics: municipalities

	Across classifications					
	Abis	**A**	**B1**	**B2+**	**B2-**	**C**
Town population	66,520	16,747	8,849	8,320	1,830	787
Median income	26,722	24,417	23,587	21,197	22,334	19,935
% in private rental housing	25.6	17.4	17.1	18.5	14.4	14.8
% moved in less than 2 years ago	12.0	11.1	11.4	11.6	9.1	8.7
% in multi-family units	80.6	40.1	26.2	26.8	11.1	6.2
Monthly rent (EUR/sqm)	20.8	15.6	12.3	10.3	10.8	8.4
House price (EUR/sqm)	5,487	3,596	2,692	1,935	1,971	1,305
Real estate sales (2014-19)	6,387	1,631	875	784	121	47
Municipalities	77	650	1,417	905	2995	30,507

Table 1.B.2 presents within-zone simple averages of socio-demographics characteristics for municipalities across the *Pinel* classification categories described in section 1.4. The data for average population, median income, the share of households in private rental housing, the share moving less than 2 years ago, and the share of the population in multi-family units, are from 2014, at the start of the *Pinel* scheme. The monthly rents are from the cross-section of 2017-2018 provided by the French Housing Secretariat. The house prices and average number of transactions are from the DV3F database over the 2014-2019 period.

Table 1.B.3: Net present value of the *Pinel* incentive

		Discount rate			
		5%	***3%***	***2%***	***1%***
Depth of rent markdown	***5%***	12.4%	13.6%	14.3%	15%
	10%	10.7%	11.7%	12.2%	12.8%
	15%	8.9%	9.7%	10.2%	10.7%

The net present value of the incentive embedded in the *Pinel* scheme is computed by assuming that the landlord receives the full benefit of the tax credit (2 percent of the purchase price each year), a 5% baseline rent-to-price ratio λ, and a nine-year commitment to affordable rents. The value is then given by $\sum_{k=1}^{9} \frac{0.02-\lambda d}{(1+r)^k}$, where d denotes the markdown of controlled rents relative to baseline market rents, and r is the assumed discount rate. See appendix 1.D.1 for a discussion of the assumptions.

Name	Period	Mechanism	Level	Affordability requirements	Increased/decreased generosity
Perissol	Jan. 1996-Jul.1999	Depreciation	80% (Y1-Y4: 10%-Y5-Y24: 2%)	None	Increased
Besson	Aug. 1999-Dec.2002	Depreciation	50% (Y1-Y5: 8%-Y6-Y9: 2.5%) option for Y10-Y15:2.5%)	Rent ceilings Resident income ceilings	Decreased
Robien	Jan. 2003-Aug.2006	Depreciation	50% (Y1-Y5: 8%-Y6-Y9: 2.5%) (option for Y10-Y15:2.5%)	Rent ceilings (close to market rate) No resident income ceilings	Increased
	Sep. 2006-Dec.2009	Depreciation	50% (Y1-Y7: 6%-Y8-Y9: 4%)	Reduced rent ceilings	Decreased
Borloo	Sep. 2006-Dec.2009	Depreciation + tax deduction	50% (Y1-Y7: 6%-Y8-Y9: 4%) +30% deduction off rental income	Rent ceilings (20% below Robien) +Resident income ceilings	Decreased
Scellier	Jan. 2009-Dec.2010	Tax credit	25%	Rent ceilings (close to market rate) No resident income ceilings 9 year commitment	Increased
	Jan. 2011-Dec.2011		22% (13% if not energy efficient)		Decreased
	Jan. 2012-Dec.2012		13% (6% if not energy efficient)		Decreased
"Social" Scellier	Jan. 2009-Dec.2010 Jan. 2011-Dec.2011 Jan. 2012-Dec.2012	Tax credit + tax deduction	Same as Scellier each year +5% per 3-year additional period +30% deduction off rental income	Rent ceilings (20% below Scellier) Resident income ceilings 9 year commitment (or 12-15y.)	Increased Decreased Decreased
Duflot	Jan. 2013-Aug. 2014	Tax credit	18% (Y1-Y9: 2%)	Rent & resident income ceilings based on "Social Scellier" levels	Unchanged/decreased
Pinel	Sep. 2014-Dec. 2021	Tax credit	12-21% (Y1-Y6: 2%) (Y7-Y9: 2%) (Y10-Y12: 1%)	Higher rent ceilings Higher resident income ceilings	Increased

Table 1.B.4: Incentive schemes for buy-to-let new residential housing in France (1996-2021)

Table 1.B.5: Eligibility exit: estimates

	(1)	(2)	(3)	(4))
Dependent Variable	New home sales	Buy-to-let (int.)	Buy to let (ext.)	Prices
3 quarters before exit	-.0428062	.0019091	.016115	-.0029532
	(.0405902)	(.0301027)	(.0182122)	(.0128662)
2 quarters before exit	.010333	.0483503	.0240921	-.0067724
	(.0435246)	(.0335806)	(.0185643)	(.0124153)
1 quarters before exit	.0660347***	.0851022***	.0447099***	.0027679
	(.0415818)	(.0319081)	(.0185127)	(.0121608)
1 quarter after exit	-.1660149***	-.1361208***	-.0423949***	-.0144896*
	(.0412847)	(0.0291161)	(.0178303)	(.0143686)
2 quarters after exit	-.2383592***	-.2275238***	-.09317***	-.0232125**
	(.0417033)	(.0295576)	(.0173303)	(.0136663)
3 quarters after exit	-.1778878***	-.1567859***	-.0678629***	-.0272029**
	(.040539)	(.028835)	(.01725)	(.0139467)
4 quarters after exit	-.218649***	-.2254706***	-.0738466***	-.0363281***
	(.0408973)	(.0296424)	(.0177864)	(.0135863)
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	90,687	90,687	90,687	32,438

Table 1.B.5 estimates the impact of losing eligibility to the incentive scheme for B2+ municipalities (relative to B1) after 2018Q4. Columns 1, 2, 3, 4 respectively plot estimates of equation 1.6 for (log) total new home sales, the (log) number of buy-to-let sales, the probability of any home sale, and the price of new units. Data are computed from the ECLN database, recording exhaustive dwelling-level information on all developer sales. Standard errors are in parentheses clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Table 1.B.6: Eligibility entry: estimates

	(1)	(2)	(3))
Dependent Variable	New home sales	Buy-to-let (int.)	Buy to let (ext.)
2 quarters before entry	-.0022547	.0059869	-.0140338
	(.0260417)	(.0164589)	(.0122359)
1 quarters before entry	.0215091	.0162246	-.001393
	(.0247059)	(.0166583)	(.0123971)
1 quarter after entry	.0415695***	.0774417***	.0296224***
	(.0262121)	(.0196763)	(.0130586)
2 quarters after entry	.0849734***	.1065969***	.0459934***
	(.0282319)	(.0210541)	(.0143595)
3 quarters after entry	.11272***	.1266174***	.0392008***
	(.0318185)	(.0231026)	(.0145558)
4 quarters after entry	.0574344**	.1202703***	.047749***
	(.032381)	(.0234635)	(.0155446)
15 quarters after entry	.2243742***	.254235***	.0924642***
	(.0406162)	(.0300801)	(.017272)
Municipality FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	134,208	134,208	134,208

Table 1.B.6 estimates the impact of gaining eligibility to the incentive scheme for B2+ municipalities (relative to B2-) after 2014Q3. Columns 1, 2, and 3 respectively plot estimates of equation 1.8 for (log) total new home sales, the (log) number of buy-to-let sales, and the probability of any home sale. Data are computed from the ECLN database, recording exhaustive dwelling-level information on all developer sales. Standard errors are in parentheses clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Table 1.B.7: Effects on the local housing stock

	Baseline effects			
"Eligibility x Post"	.0116819***	.0159897***	.0096321**	.0233584***
	(.0025491)	(.0026204)	(.004432)	(.0063502)
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Time-varying controls	No	Yes	Yes	Yes
Weighting by pop.	No	No	Yes	Yes
HS elasticity	No	No	No	Yes
Observations	42,631	40,598	40,598	40,349
	Heterogeneous effects			
"Eligibility x Post"	.037374***	.0423141***	.0222928	.0346503**
	(.0113679)	(.0110031)	(.0148743)	(.0147433)
"Eligibility x Post x S_c"	-.0507972**	-.0522198**	-.0250971	-.0226713
	(.0222711)	(.0213539)	(.0358387)	(.0361204)
Municipality FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Time-varying controls	No	Yes	Yes	Yes
Weighting by pop.	No	No	Yes	Yes
HS elasticity	No	No	No	Yes
Observations	42,094	40,114	40,114	40,114

Regressions coefficients are shown for the estimation of equation 1.9 in the sample of B2 municipalities. The "Eligibility x Post" coefficient indicates the DiD estimate of β_R, the interaction of post-entry date and eligibility (membership of the B2+ group). The "Eligibility x Post x S_c" coefficient is the estimate of θ_R, the heterogeneous impact of eligibility on the local housing stock by the extent of $S_c = \sum_i \frac{H_{i\to j(c)}}{H_{j(c)}} s_{ij(c)}$, the model-implied measure of the concentration of ownership for the commuting zone j to which c belongs. The model predicts that a more dispersed ownership (a lower S_c) entails a stronger effect of subsidies on housing production. "Time-varying controls" include the interaction of year fixed-effects with initial characteristics of the municipality (median income and population). "Land share" indicates the presence of controls for the interaction of "Eligibility x Post" with the baseline urbanized land share in the municipality. Standard errors are in parentheses clustered at the municipality level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

1.C Data construction for Chapter 1

I use data drawn from a wide array of administrative sources, including exhaustive individual data on the universe of building licenses requests, housing starts, land purchases, housing transactions, developer residential sales, and social housing units in France. A confidential and exhaustive panel of all taxpaying households in France allows me to track the individual determinants and consequences of investors' decisions to invest in the housing sector. Finally, I also rely on municipality-level measures of aggregate outcomes, spatial and housing mobility, as well as municipality-industry measures of employment, employee earnings, and firm creation for all c. 35,000 municipalities in France, in order to assess the consequences of place-targeted housing policies at the local level. This appendix provides additional detailed information on the construction of the data, sample selection, and data transformation performed.

Housing markets

Home sales To study the impact of the policy on local markets for both newly built units and existing homes, and track its role for local home prices, housing market activity and land values, I use two main sources of administrative data on home sales. First, I avail myself of exhaustive and never-used before data recording all housing and land transactions in France from January 2010 to December 2019, recorded by the French Treasury for tax purposes in the *DV3F* database. For home sales, I restrict the sample to sales of entire houses and apartments, which represent around 850,000 annual housing transactions over the sample period. Transaction details include the name of the municipality, the exact geo-coding of the address, and, for most transactions, the price and exact square footage of the unit sold. Using specific legal codes defining the VAT and transfer tax treatment of the sale, I am able to determine whether a sale corresponds to a new or an existing home, and whether the sale was an "off-plan" sale (*Vente en l'Etat Futur d'Achevement* or VEFA) - the main mechanism through which developers sell rental property in France.

Second, I use a restricted-access exhaustive survey of developers, the *Enquete sur la Commercialisation des Logements Neufs* (ECLN), which precisely records detailed information on all new housing sales performed by developers at the unit level. Relative to the DV3F data, the

ECLN includes additional information on the characteristics of the unit, the identity of the developer, and, for most units since the 2010-2011 wave, the final destination of the purchase (owner-occupied housing, subsidized or unsubsidized rental property, social housing...), as well as its exact location. The survey is exhaustive and of high quality, with projects being followed over time for several quarters; it is however restricted to housing projects comprised of five or more units, thus including all large multi-family projects but potentially missing small individual constructions still eligible under the Pinel scheme.

Land transactions I avail myself of two main sources for land sales and prices in France. First, I exploit the *DV3F* database but restrict transactions to sales of constructible land ("*terrains a batir"*). Second, I use the *Enquete sur le prix des terrains a batir*, an exhaustive survey recording the purchase price, final use, and overall area of all constructible land sold in France every year, to examine the consequence of the policy for land prices.

New construction To evaluate the impact of the *Pinel* scheme on house-building, I use administrative confidential data from the *Sit@del2* database combined by the French Housing authority. The Sit@del2 data contain exhaustive information on the universe of housing permits requested and include details about the exact date of the request, its approval date and current status, the geographic location, a summary description of the project, and the number and type of units built, for each housing project. The database allows me to aggregate all housing authorizations and housing starts (separately for single- and multi-family units), at the dwelling level across all municipalities in France.

Housing stock I rely on exhaustive annual data from the cadaster - the so-called *Fichiers Fonciers*, a tax-based registry of the universe of housing units in France, which records extensive information on all housing units at yearly frequency, as of January 1st. I aggregate data at the municipal level, in order to compare a variety of relevant outcomes between eligible and uneligible locations, including the total number of local dwellings and the share of rental and multi-family units. As a complement, I also use data from INSEE, the French statistical institute, which combines the annual Census with other sources (including INSEE's regular Housing survey, housing tax files, and other information) to produce information on all housing units – occupied or not – in France. The resulting *Base Logements*, available at yearly frequency until 2017, provides information on the quantity, use, and nature of all housing units in a municipality, including detailed information on tenure status, available amenities, turnover rates, and ownership status.

Social housing To study the extent of crowding-in or crowding-out of social housing by intermediate housing construction incentivized through the Pinel scheme, I use the *Répertoire des logements locatifs des bailleurs sociaux*, an exhaustive and legally mandated repository of all social housing units in France. The RPLS provides an annual assessment of the number of social housing units (occupied or vacant) in the municipality, starting in 2012. Since reporting to the RPLS is mandatory for all owners of social housing, its data provides a highly reliable source of information on the extent, exact location, and characteristics of new social housing supply, at yearly frequency, for all municipalities in France.

Investor behavior

Tax returns I use the universe of French individual tax returns from 2006 to 2019, an exhaustive sample of more than 30 million tax units each year (collectively known as *Fichiers POTE - panelisables*). The data, the panel structure of which allows me to track taxable households over time, include detailed information on household characteristics, all forms of taxable capital, pension, and labor income, any tax deductions and shelters used, the amount (if any) invested under subsidized housing schemes. In order to measure landlord status, I define as landlords all households who file a so-called *2044* schedule, and register either a strictly positive amount of income from rental housing, or deduct a strictly positive amount of rental losses ("negative gearing"). I merge individual tax returns data to mandatory additional schedules filled by initial investors in the subsidized rental investment scheme (known as schedule *2044-EB*, for *engagement bailleur*, i.e. "landlord commitment"): for almost all buyers under the *Pinel* scheme who fill an online tax return in 2016, these data record the exact location of the property they purchased, the price paid, the square footage of the unit, as well as additional information on the tenant and unit characteristics.

Housing stock ownership I make use of the FILOCOM/FIDELI database, an exhaustive repository of all c. 37 million housing units (as of 2019) in France, available every other year (and every year starting after 2018). Data on individual housing units obtained from the property tax collection process are merged by tax authorities to summary information from income tax returns on resident households. In addition to several pieces of information on the dwelling, including the income and household composition of its residents, its construction year, or the last year in which it was sold or donated, the database records the municipality of the address to which the property tax (the *taxe fonciere*) is sent. For each non-owner-occupied property, this allows me to observe the residence location of the rental investor for each rental dwelling, a key input in my gravity approach to housing investment. I match these data to a matrix of great-circle, road, and cultural distance between the commuting zones or provinces to which the municipalities of the owner and the renter belong, in order to examine the role of spatial frictions in residential investment behavior.

Wealth and assets Since tax returns data do not record exhaustive information on overall assets and wealth (see e.g. Garbinti, Goupille-Lebret, and Piketty (2021) for an attempt at reconciling tax and survey data on wealth in France), I also exploit the *Enquete Patrimoine* (*EP*), a French equivalent to the U.S. Survey of Consumer Finances, which records detailed information on household income, assets, and portfolio composition. The survey is a short panel which interviews household twice over a four-year period (I use the 2018 vintage). It is matched to summary information on income and family status from household tax returns in the past year, and over-samples wealthier tax units. I use restricted-access microdata from the *EP* in order to measure landlordship and its correlates at the individual level. The data allow me to observe, for each surveyed individual, the level and allocation of their wealth across various asset types, including the value and nature of housing assets owned by a household. Landlords are defined as individuals who own at least one housing unit (in addition to their primary residence) that they lease for the entire year: these correspond to c. 10 percent of French households (c. 3 million).

Data on the local economy

Population I use data on municipal population assembled by INSEE through the annual Census, to estimate the extent of population mobility induced by the scheme towards targeted areas. In addition, I employ aggregate data on municipality-to-municipality annual migration flows from the French statistical institute INSEE, to measure the extent to which the policy induced population mobility towards targeted areas, and away from untargeted regions. To study the composition of aggregate flows and assess how the socio-demographic composition of renters and movers, I use individual data from the Census on mobility decisions from INSEE's *Fichiers details - Mobilite residentielle*.

Renters and low-income households To look at specific mobility effects among lower-income households who rent, I rely on data from the *Caisse d'Allocations Familiales*, France's main outlet for social benefit payments. In particular, the CAF records the annual number of beneficiaries of each of its three main housing allowances, the APL (*Allocation Personnalisee au Logement*), ALS (*Allocation de Logement Sociale*), and ALF (*Allocation de Logement Familiale*), in each municipality, through the *Fileas BCA* database. These allowances are designed to help mostly renters, but also homeowners, cover their regular housing expenditures.[72] I use the *Fileas BCA* data to obtain counts of beneficiaries of any housing allowance, as well as separate counts of renters and homeowner beneficiaries. I also use the CAF files to measure the percentage of poor households in the local population, defined as those earning less than sixty percent of median income per consumption unit.

Local economic activity To test whether the policy affects local economic activity, I use a series of administrative databases on municipality-level outcomes. I first exploit establishment-level data from the *Repertoire des Entreprises et des Etablissements* database,[73] comprised of information on the number of employees, detailed industry classification, and exact location for each establishment in France, which I aggregate at the municipality-year level, to estimate the annual stock of firms, as well as entry and exit rates, in each industry and municipality, before and after the implementation of the policy. I also use data from the French

[72] Bozio et al. (2015) provide a detailed overview of France's various personal housing benefits.
[73] Base SIRENE (BPE) - 2007-2019, INSEE [producteur], ADISP [diffuseur].

Social Security system (*Unions de Recouvrement des cotisations de Sécurité Sociale et d'Allocations Familiales* or URSSAF), aggregated at the municipality-industry-year level, on the total number and overall wages of salaried workers by 4-digits industry, to study the impact of the policy on employment and wage rates in the real estate and construction sectors.

Land use The impact of subsidizing housing supply on urban sprawl and expansion requires granular, geo-localized information on land use and land cover. I avail myself of the *Observatoire de l'artificialisation des sols* database, which combines satellite information with administrative records, and constitutes the most precise and exhaustive repository of information on land use at a highly granular definition (at the level of *"parcelles"*, or cadastral plots). The data allow me to compute the amount of land converted each year from "natural" use to residential or industrial use, within each French municipality. I also use the *Corine Land Cover* database, a satellite measure of land use, to assess the share of urbanized land in each town.

Institutional arrangements

Subsidized housing scheme eligibility I exploit data from the French Housing Ministry detailing the eligibility status and zoning rules applicable to each municipality in France. The data record the zone applicable to a municipality ("A","Abis","B1","B2", or "C"), both from 2009 to the September 2014 change and afterwards. The law determines the eligibility of an area to the subsidized investment scheme, and, for eligible locations, both the rent ceiling per square meter applicable and the maximum income of the residents. I also gathered data on all individual eligibility agreements granted to B2 or C municipalities under the Pinel scheme. I obtained restricted-access data from the Ministry of Defence on the coverage of towns by "garrison towns" revitalization contractual agreements (*Contrats de Revitalisation des Sites de Defence*), since the 2018 reform expanded Pinel eligibility to cover towns that had lost a military unit and were covered by such a revitalization contract (see 1.D.3). Finally, to study the diverging dynamics of neighbouring towns when they differ by eligibility status, I also employ an adjacency matrix of all municipalities provided by OpenStreetMap, an open-source software recording adjacency links in the graph of all c. 35,000 French municipalities.

1.D Institutional background

1.D.1 Subsidy schemes to new housing in France

Rental property policies in France France spends substantial amounts of budget support for housing policies. As of 2019, the country dedicated c. €40 billion a year (around $45 billion, close to two percent of French GDP) to various housing market demand and supply-side subsidies. These include around EUR 22 billion in direct rental assistance to low- and middle-income tenants, c. EUR 3 billion in interest rate subsidies for social housing construction and exploitation, and interest-free loans to middle-income owners; c. EUR 2 billion in building subsidies, and around EUR 13 billion in a variety of fiscal expenditures, from reduced corporate tax and VAT rates for social housing construction, to tax deductions for environmental upgrades of primary residences, to tax credits associated with rental property investment (see figure 1.D.1).

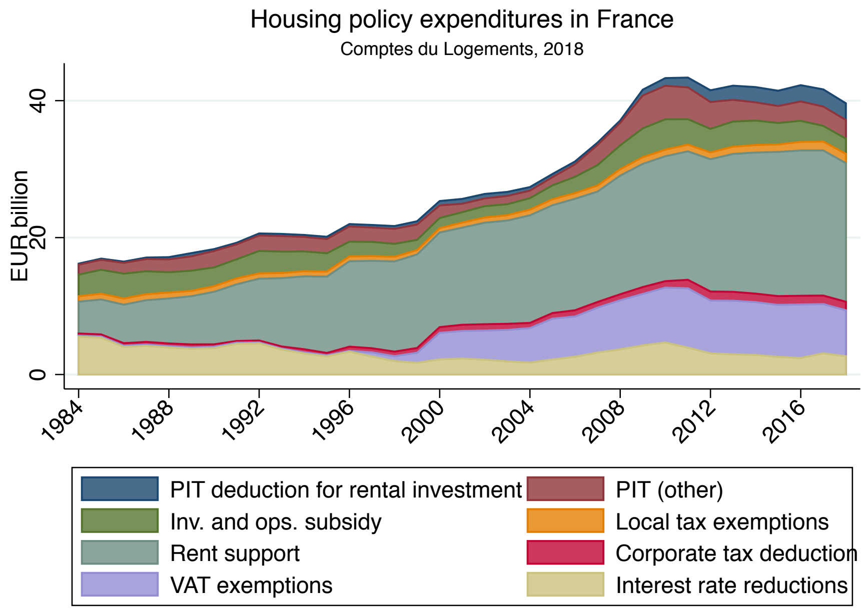

Figure 1.D.1: **French housing policy costs**

Various housing policy expenditures in France were estimated to cost close to 40 billion €as of 2018, or c. 2 percent of French GDP. They include personal income tax deductions for rental investment, the focus of this paper, as well as a variety of other tax expenditures, VAT/corporate tax/local taxation exemptions, interest rate subsidies, operations and investment subsidies to housing developers, or direct cash allowances to tenants. Figures are from *Comptes du Logement*, an annual report compiled by the French Housing Secretariat on the total costs of policies supporting housing. Sources: *Service de la donnée et des études statistiques, Sous-direction des Statistiques du Logement et de la Construction.*

Tax credits for rental property investment The latter are the focus of the empirical exercise in this paper. They allow for either accelerated depreciation deductions off the investor's personal income tax basis,[74] or direct reductions in the investor's tax liability. A succession of tax schemes were put in place starting in the 1990s, with varying deduction rates and limitations related to affordability guidance, targeted areas, rent ceilings, and tenant income conditions. They were renamed after each Housing Minister who introduced changes. Table 1.B.4 provides summary information on the schemes, which are describe in more detail below. [75]

[74]In France, depreciation is generally not allowed as a deductible expense for owners of rental property, unless the unit is furnished and rented under a specific tax status, the *Loueur en meuble non professionnel* for a short period of time.

[75]Successively: Perissol, Besson, Scellier, Robien, Duflot, and Pinel. Historian and sociologist Vergriete (2013) provides a long-run view of the political implications of these policies.

- From 1996 to August 1999, the *Perissol* scheme was in place. It allowed households to depreciate 80 percent of the purchase cost of the rental property (10 percent of the investment in the initial four years, then another 2 percent a year for 20 years, with no affordability requirement). The depreciation could be used to reduce taxable rental income, and any excess could offset overall taxable income up to a limit of c. EUR 15,000 per year.

- From 1999 to 2002, the *Besson* law reduced the generosity of the accelerated depreciation mechanism. It allowed a depreciation of only 50 percent of the cost of newly built residential assets to offset the individual income tax liability of the investor over nine years (8 percent in the first five years, then 2.5 percent a year for four years), with an option for an additional 2.5 percent a year depreciation for 6 years, pushing the total potential for tax depreciation to 65 percent over 15 years. The cap on excess depreciation (negative gearing against non-rental taxable income) was lowered to c. EUR 10,000 annually. Moreover, below-market rent ceilings and maximum resident income conditions were put in place, and made dependent on whether the investment occurred in Paris, in large cities, or elsewhere, with a view to turn the scheme into a form of inclusionary zoning.

- From 2003 to 2006, the more generous *Robien* law kept the same depreciation conditions as the Besson law, but raised rent ceilings by 10 to 50 percent (depending on the location) and abolished the maximum resident income provision, therefore increasing net incentives for taxable households to invest in individually-owned rental housing.

- From 2006 to 2009, the so-called "re-focused" *Robien* law slowed down the pace of allowed depreciation (to 6 percent a year for seven years, and four percent for an additional two years), removed the option to extend it for another six years, and lowered rent ceilings in low-density locations. An alternative scheme, the *Borloo* law, allowed households to add to the basic *Robien* tax shelter an option for a 30 percent deduction off rental income, under the condition that rents be capped c. 20 percent below *Robien* rent ceilings.

- From 2009 to 2010, with a view to kickstart the recovery after the Great Financial Crisis, the newly introduced *Scellier* scheme reshuffled tax subsidies to owner-operated rental housing. It changed the definition of eligible areas and excluded small towns from the subsidy scheme. It also replaced the bonus depreciation deduction system by a non-refundable fixed-rate investment tax credit of 25 percent of the purchase price (5 percent in the first year, then 2.5 percent a year for eight years), which directly reduced an investor's tax liability instead of offsetting taxable income. An alternative scheme, the *"social" Scellier* law, similar to the earlier *Borloo* scheme, allowed households to add to the basic *Scellier* tax benefits an option for a 30 percent deduction off rental income, under the condition that rents be capped c. 20 percent below *Scellier* rent ceilings. Due to budget constraints amid the euro area financial crisis, subsidies to newly produced rental housing investment under the *Scellier* and *social Scellier* schemes were first sharply reduced in 2011: the rate of the investment tax credit was lowered to 22 percent for units respecting strict environmental and energy efficiency, and to 11 percent for other units. The tax credit rate was then drastically cut down in 2012 to 13 percent for energy-efficient units, and to 6 percent for other units. Maximum purchase prices per square meter, which varied by zone, were also implemented starting in 2012.

- From January 2013 to September 2014, the Duflot scheme allowed for an 18 percent tax credit (2% annually over nine years), but implemented stringent rent conditions, and excluded B2 and C cities (unless they obtained an agreement from the prefect) from the benefit of the scheme.

- Finally, after September 2014, the *Pinel* scheme reshuffled eligible areas, upgrading towns with c. 10 million inhabitants into zones with higher rent and income ceilings, and allowed for a more flexible rental commitment duration (from 6 to 12 years) and the possibility to rent a unit to one's relatives. More information on the *Pinel* scheme is provided in the main text. Figure 1.D.2 provides information derived from individual tax data on the take-up of the Pinel scheme, distinguishing purchases by the year in which taxpayers first claim the rebate (completion year) and by the delay since the purchase of the unit.

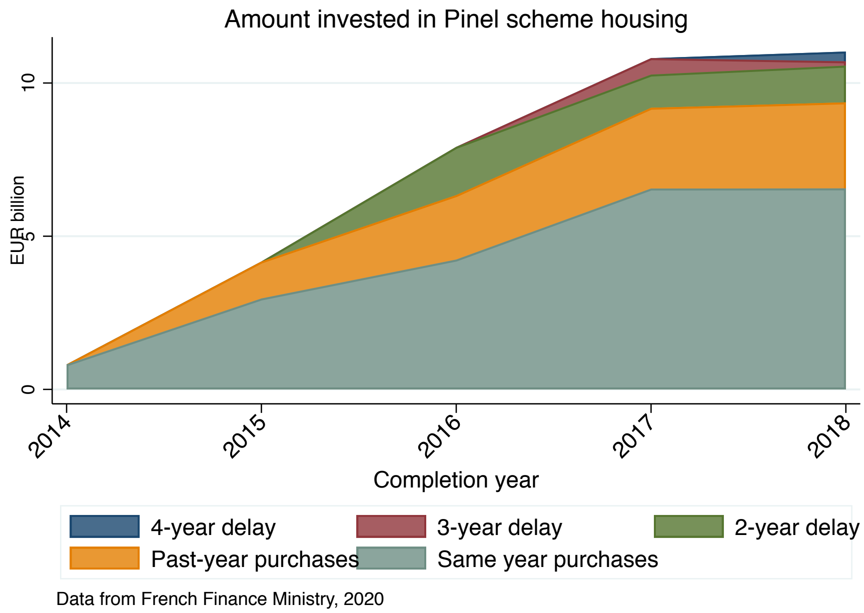

Figure 1.D.2: **Total value of Pinel take-up by completion year**

The figure plots the annual value of Pinel investments by the year they were first reported by tax filers – corresponding to the year the unit was first rented, along with the delay between the reporting year and the year the purchase was effectively made. For example, Pinel investments first reported for tax filing year 2017, but for which the purchase occurred in 2015, represent c. EUR 1.08 billion, while Pinel investments corresponding to purchases made in 2017 and rented for the first time in the same year were c. EUR 6.6bn. Sources: *Annuaire statistique* database, the annual report of the French Treasury reporting all personal income tax-related information - see appendix 1.C for a detailed description of the data.

Public interest in landlord incentives Figure 1.D.3 plots the evolution of interest in each of the last four schemes in operation (Robien, Scellier, Duflot, and Pinel) since 2004. Interest for each of the schemes (as measured by GoogleTrends search volume, normalized to 100 at the maximum value) follows its implementation and timing of availability. Within the period in which schemes are available, interest gauged by Google searches spikes towards the end of each fiscal year, and close to planned eligibility deadlines, demonstrating the salience of tax motives, and consistent with empirical evidence on the actual timing of purchases which tend to bunch towards the end of fiscal years and immediately before the end of eligibility.

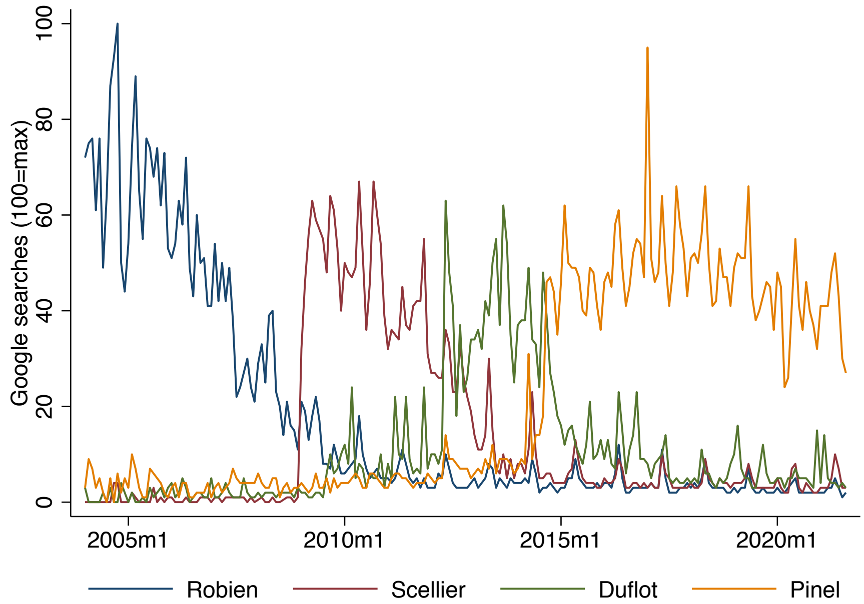

Figure 1.D.3: **Interest in tax-favored schemes**

The figure plots the Google Trends index of search intensity (normalized to a maximum value of 100) for the last four schemes in operation (Robien, Scellier, Duflot, and Pinel). The successive tax regimes are described in table 1.B.4.

Net present value of the *Pinel* incentive Rents under the Pinel scheme are theoretically subject to a ceiling amount, designed to target a 10 to 20 percent discount relative to average market rents in a zone. In 2017, it was set at C=EUR 16.83/sq.m. in area A bis; EUR 12.5/sq.m. in A; EUR 10.07/sq.m. in B1; EUR 8.75/sq.m. in B2. For example, a *Pinel*-eligible 50 sq.m. apartment in Marseille (zone A) would be subjected to a $C \times M = 12.5 \times (0.7 + 19/50) = 13.5$ EUR/sq.m. ceiling, while a *Pinel*-eligible 30 sq.m. apartment in Toulouse (zone B1) would face a $C \times M = 10.07 \times 1.2 = 12.1$ EUR/sq.m. ceiling. Resident household income was capped (for a couple with one child) at 72,737 €in area Abis; 66,699 €in A; 48,596 €in B1; and 43,737 €in B2. Area-specific income ceilings vary with family size, with larger households subject to looser maximum income guidelines. The income ceiling only applies in the first year of the lease. Nationwide, the median household income for a couple with one child was 43,880 €in 2017. Income ceilings cover close to 80 percent of the population and rarely bind.

The effective value of the *Pinel* subsidy to new investment varies across locations depending on the depth of required rent discounts relative to market rents. Nonetheless, using data on effective market rents and ceilings applicable under the Pinel scheme, it is possible to estimate a range of possible values for the effective net present value of the incentive embedded in the Pinel scheme. I make data-driven but highly conservative assumptions for the discount rate, the extent to which affordability requirements fall below market rents, the duration of investors' commitment to affordability, and average market rent-to-price ratios. I take a baseline rent-price ratio of 5 percent, close to the median in the sample of B2+ municipalities. During the period of interest, interest rates on 10-year French government bonds (the relevant discount rate for government-guaranteed future transfers) were in the range of 1 to 2 percent; this value is similar to the interest rates on new mortgages documented by the Bank of France over the period. Among B2+ municipalities, the mean rent markdown implied by the *Pinel* regulation relative to baseline market rents was 6 percent, and the 90^{th} percentile was 17 percent (see figure 1.D.4), so that central estimates of the NPV of the incentive should cluster around values around 5 to 15 percent for the depth of the rent markdown. Lower baseline rent-price ratios, a 12-year commitment, or non-binding rent markdowns would deliver even larger values of the NPV of the incentive. Table 1.B.3 provides a sensitivity analysis for the value of the tax incentive net of affordability requirements. While preferred estimates are in the upper-right quadrant of the table, even for assumptions of very high discount rates and a severe bite of rent discounts, the net value of the incentive is still above or close to 8-9 percent of the initial purchase price of the asset.

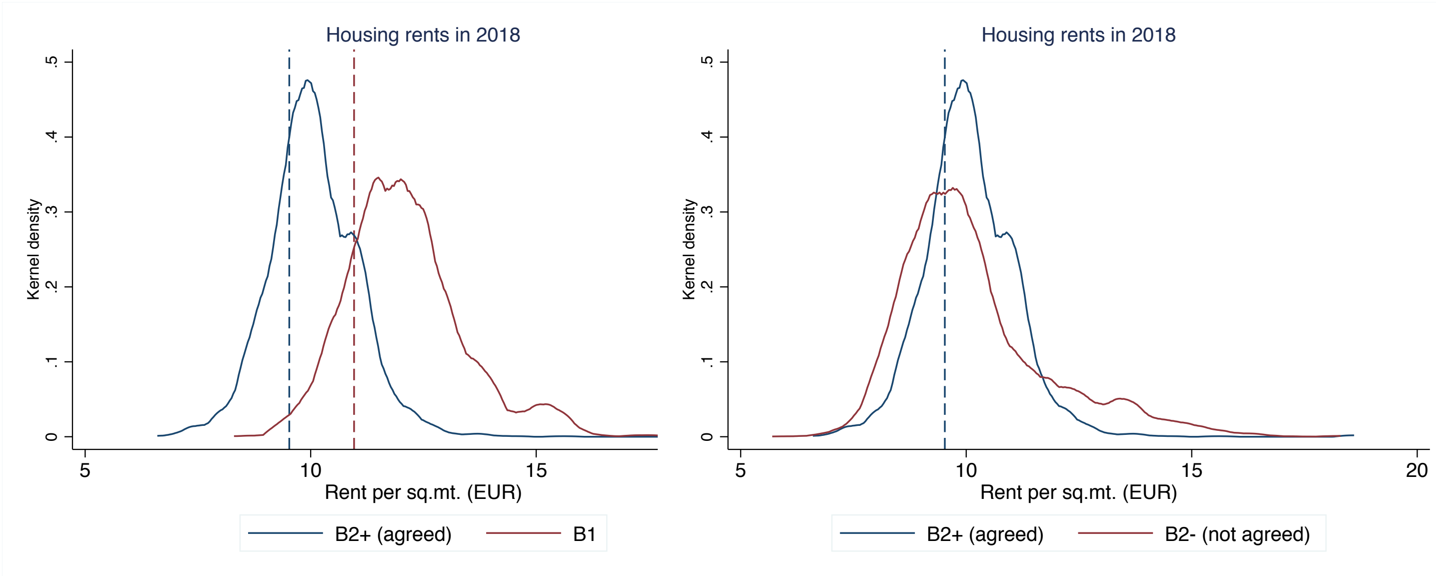

(a) B2+ and B1

(b) B2+ and B2-

Figure 1.D.4: **Rents by zoning area**

The figure plots the rents per square meters for apartments in 2017-2018 for cities in the main treatment group – B2+ municipalities with an agreement (blue line), and their counterparts used as controls – always-treated B1 and never-treated B2- municipalities. Vertical dashed lines mark the rent ceiling (for a 50 sqm apartment) under the *Pinel* scheme for each area's corresponding color. Rents display substantial heterogeneity within each area, and substantial overlap across areas. Sources: French Housing Secretariat – see appendix 1.C for a detailed description of the data.

1.D.2 Legal ranking of municipalities by property market tension

Higher housing price markets were broadly allocated to higher-ranked areas. One may think of the classification as a slightly more granular version of the United States' "Difficult Development Area" status, whereby areas with higher housing costs are ranked at a higher degree of priority for government support to low-income housing. Nonetheless, there was both substantial heterogeneity in housing demand and supply conditions *within* each one of the five areas, and significant overlap *across* them. As an example, figure 1.D.5, panel (a) compares the (kernel density smoothed) distribution of municipality-level median home prices from 2014 to 2019 for municipalities in the B1 and B2+ areas; panel (b) of the same figure compares agreed municipalities in the B2 zone with B2 municipalities without an agreement (*B2-* towns). The distribution of median house prices by municipality closely track each other in agreed and non-agreed B2 municipalities, although there is substantial overlap between the agreed B2 and B1 municipalities. Figure 1.D.6 does the same comparisons for the share of urbanized land. This time, the distribution of urbanized land cover by municipality more closely resemble each other in B2+ and B1 municipalities, although there is also substantial overlap between the B2+ and B2- cities.

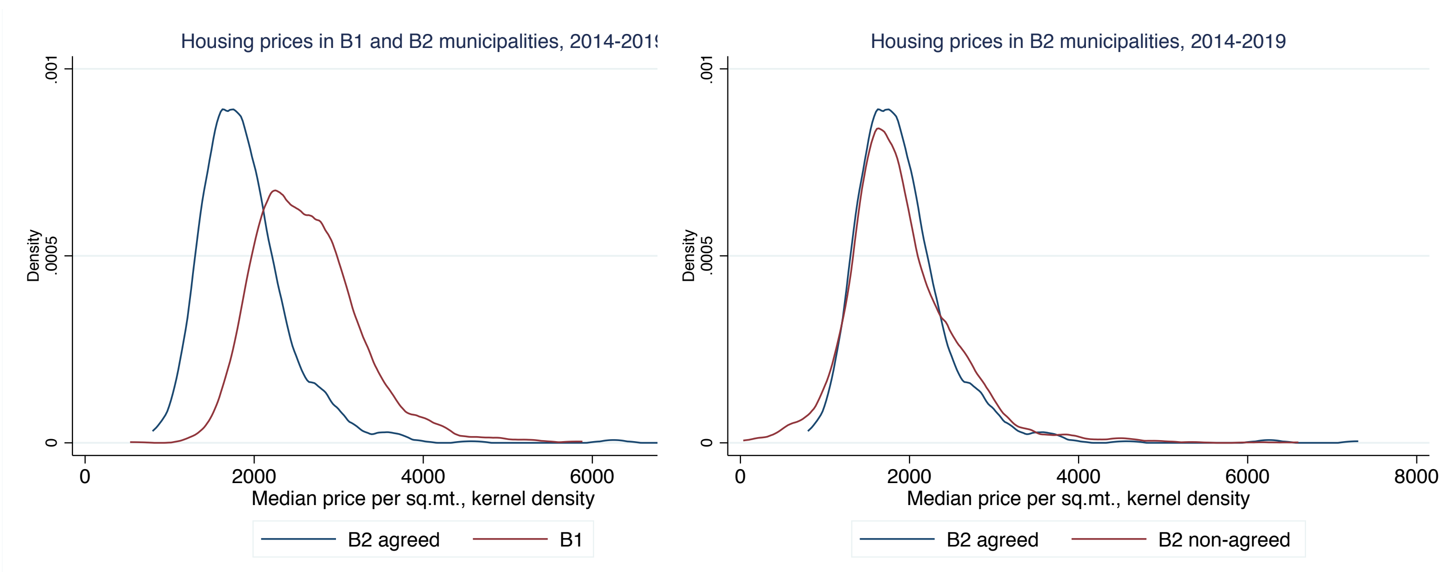

(a) B2+ and B1

(b) B2+ and B2-

Figure 1.D.5: **Median house prices by zoning area**

The figure plots the median home price for apartments and houses sold from 2014 to 2019 available in the DVF database for cities in the main treatment group – B2 municipalities with an agreement (blue line), and their counterparts used as controls – either B1 (panel (a)) or B2 (panel(b)) municipalities (red line). Home prices display substantial heterogeneity within each area, and substantial overlap across areas. Sources: *DVF* database, an exhaustive record of housing transactions in France made available by the French Finance Ministry. Sales restricted to houses and flats with an available transaction value and square footage - see appendix 1.C for a detailed description of the data.

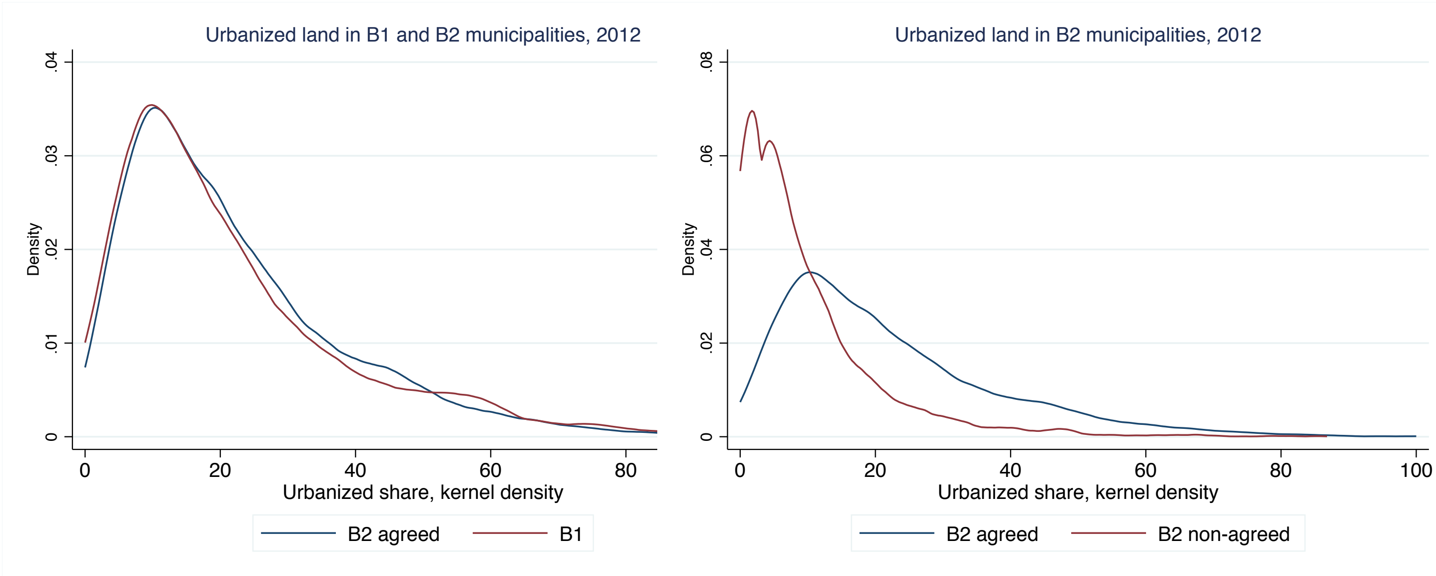

(a) Agreed B2 and B1

(b) Agreed B2 and non-agreed B2

Figure 1.D.6: **Share of urbanized land by zoning area**

The figure plots the share of urbanized land available in the Corine Land Cover database for cities in the main treatment group – B2 municipalities with an agreement (blue line), and their counterparts used as controls – either B1 (panel (a)) or B2 (panel(b)) municipalities (red line). Urbanized shares display substantial heterogeneity within each area, and substantial overlap across areas. Sources: *Corine Land Cover 2012* database, a land cover inventory initiated in 1985, providing detailed information on the use of local land (land cover) broken down into 44 classes. Data are made available at the municipality level by the French Environment secretariat.

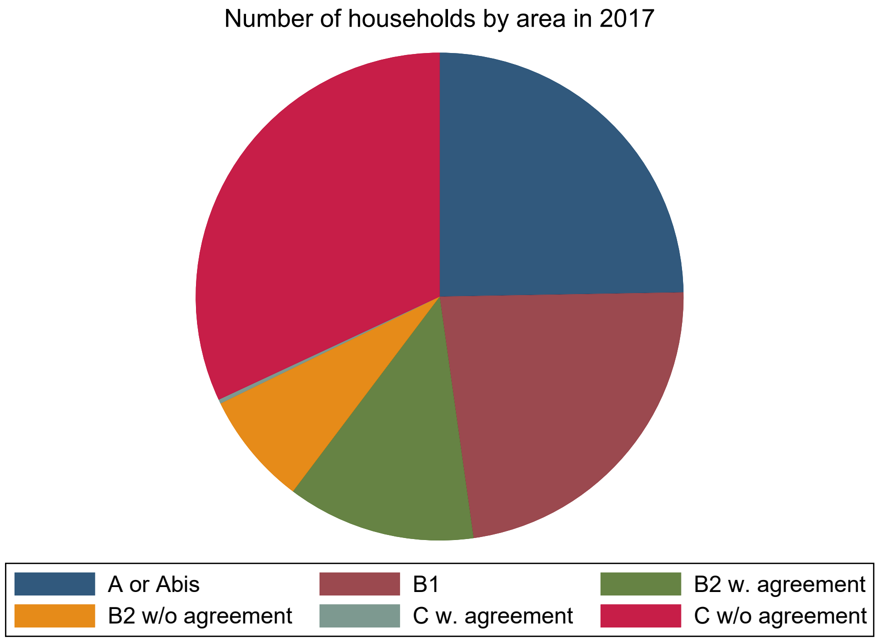

Figure 1.D.7: **Population by *Pinel* zoning**

This figures provides the breakdown of household numbers (as of 2017) by municipality type. Municipalities were mapped into one of five areas, determining both their eligibility to the *Pinel* scheme, and rent guidance and income restrictions. Areas are ranked in descending order of estimated property market "imbalances", with Abis being the most highly-demanded areas and C being the least demanded. B2 and C areas are further broken down by agreement status (B2+/C+ are B2 or C towns with a prefect's agreement that are eligible to the *Pinel* scheme).

1.D.3 The "garrison towns" special treatment

The *Pinel* policy was discontinued for B2 municipalities with an agreement after January 1, 2018 for new projects (March 2019 for sales). However, the same legislation expanded the policy coverage to include all municipalities that were (or had been part at some point within the last 8 years) part of a program to revitalize former defence garrison towns which had lost military units due to a reorganization of French forces. These towns were part of contractual agreements between the Ministry of Defence and local authorities labelled CRSD or *Contrats de Revitalisation des Sites de Defense*. After January 2018, any town that had been recently covered by a CRSD plan was eligible to the *Pinel* scheme, independently of its classification. Therefore, CRSD-covered towns in C or B2 zones without an agreement prior to 2018 gained eligibility, while CRSD-covered B2+ towns kept their eligibility to the *Pinel* scheme in spite of the 2018 reform. For A-Abis-B1 towns, which all kept their eligibility to the scheme after

2018, the CRSD expansion did not affect the availability of *Pinel* scheme investments.

1.E Stylized facts: additional evidence

1.E.1 Rental housing and mobility

Labor mobility and private rental housing are closely associated I document that residential mobility and private rental housing are tightly linked, suggesting a spatial equilibrium framework should focus on the provision of rental, rather than owner-occupied, housing.[76] Using Census data aggregated at the municipality level, panel (a) of figure 1.E.1 evidences that, in the (population-weighted) cross-section of France's 35,000 towns in 2018, a larger share of private rental properties in the local housing stock of a municipality is tightly correlated with higher housing turnover (as measured by the share of units which households moved in less than two years earlier). On average, a one standard deviation (10 percentage points) increase in the share of private rental dwellings in the local housing stock is associated with a 0.7 standard deviation (2.5 p.p.) higher share of households who moved in recently.[77] Using individual-level information on the distribution of tenure duration by status (renter versus owner-occupier) among mainland France's more than 29 million primary residences in 2019, panel (b) of figure 1.E.1 shows that occupancy in the private rental sector is substantially shorter and left-skewed, relative to owner-occupied housing.[78]

While turnover at the dwelling level is relevant to housing market fluidity, some of the differential mobility of renters could be driven by a higher frequency of within-town moves. Using detailed town-to-town mobility data from the Census, panel (c) of figure 1.E.1 exam-

[76] Spatial equilibrium models often abstract from differences in tenure status, but Blouri, Büchler, and Schöni (2021) and Favilukis and Van Nieuwerburgh (2017) both characterize residential tenure mode as an endogenous household decision. Fixed transition costs in and out of owner-occupied housing can be large, making renters more prone to move. Monetary fixed costs specific to homeowner mobility include real estate transfer taxes, realtor fees, or capital gains taxes after a nominal price appreciation (see e.g. Levy (2021) for quasi-experimental evidence of the impact of the latter in France). Other, non-pecuniary, differential fixed costs may include deeper local ties for homeowners or behavioral biases such as endowment effects. The higher mobility of renters documented in this section could thus be due either to a *causal* effect of tenure (e.g. through differential mobility costs), or *selection* patterns, as households more prone to move due to observable and unobservable characteristics tend to choose renting rather than owning (Oswald, 2019).

[77] On the other hand, the correlation of this measure of household mobility with the share of *public* rental housing is almost nil (panel (1.G.1a) of appendix figure 1.G.1). Public housing (provided by the government at deep discounts against market rents) represents close to 40 percent of all rental units in France, or close to five million units.

[78] The median tenure for renters in the private sector is 2 years [interquartile range: (0,5)], while the median duration for owner-occupied units is 10 years [interquartile range: (4,20)], and 5 years [interquartile range: (2,12)] in *public* rental housing.

ines longer-range geographic mobility, which is more relevant to the spatial reallocation of labor. It displays the close association between the degree of inwards migration in a town and local private rental dwellings availability.[79] In addition, not only do towns with a higher proportion of private rentals receive more movers, but inwards migration there comes from further away. Weighting bilateral inwards moves in a municipality by the distance between the previous residence and the current town, panel (d) of figure 1.E.1 shows that towns with more private rental dwellings receive longer-distance inwards movers. Therefore, both the extensive (share of movers) and intensive (average distance among movers) margins of spatial mobility closely co-move with the availability of private rental housing in a location. These regularities motivate a model where "grounded" owner-occupiers are relatively immobile across space, but "footloose" renters move in response to - and partially arbitrage - utility differentials across locations, making the availability of *rental* housing a key determinant of overall mobility.[80]

[79] Panels (b) and (c) of appendix figure 1.G.1 show no such close linkage exists between cross-city moves and the share of *public* rentals in the local housing stock, as both tend to be *negatively* correlated. The same regularity holds for outwards migration, which is strongly and tightly positively correlated with the availability of private - but not public - rental housing.

[80] The close linkage between rental housing and mobility is not specific to the French context. In the United States, according to the Census Bureau for the year 2017, renters moved at an annualized frequency of 21.7 percent, against only 5.5 percent for owner-occupiers.

(a) Private rental housing and housing turnover

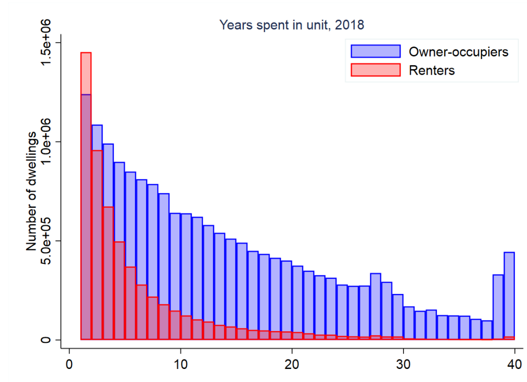

(b) Tenure duration by occupancy status

(c) Private rental housing and inwards mobility

(d) Long-distance moves and private rentals

Figure 1.E.1: **Private rental housing and spatial mobility**

Panel a plots the (population-weighted) relationship between the share of the population living in private rental units in a municipality, and the share of the population living in units which they moved in less than two years earlier, in 2018. Panel b plots the distribution of tenure duration (measured as the number of years since the household moved into the property) separately for renters in the private sector and for owner-occupiers. Same-year movers are not shown, and the maximum tenure duration is censored at 40 years for legibility purposes. Panel c plots the (population-weighted) relationship between the share of the population living in private rental units in a municipality, and the share of the population which moved into the municipality in the last year, in 2018. Panel d plots the (flow-size weighted) average distance of moves ($\bar{d}_i = \sum_j \frac{F_{ij}}{\sum_j F_{ij}} d_{ij}$) by inwards movers from j to to municipality i in the last year (in logs), against the share s_i of the population living in private rental units in i, in 2018. The distance of moves is computed as the haversine distance between the centroids of the municipality of origin and the destination municipality. Panels a, c, and d present binned scatter plots where the average value of the y-axis variable is plotted against each (population-weighted) centile of the x-axis variable, and data for these panels come from France's national statistical institute *Base Logements*, a housing survey matched to the Census and combined with register data. Data for panel b are computed from register data (the *FIDELI* database) on the universe of c. 29 million dwellings used as primary residences in France.

The differences in spatial and residential mobility behavior between private sector renters and owner-occupiers correlate with distinct observable characteristics of agents choosing each of the two tenure modes. Table 1.E.1 provides demographic information on the refer-

Table 1.E.1: Characteristics of household by tenure mode

	Owner-occupiers	Private sector renters	Public sector renters
Household size	2.30 (1.21)	1.89 (1.17)	2.30 (1.43)
Age of reference individual	58.7 (16.5)	43.6 (18.2)	51.3 (16.8)
Full-time employed	0.56 (0.50)	0.61 (0.49)	0.52 (0.50)
Out of labor force or retired	0.40 (0.49)	0.20 (0.40)	0.31 (0.46)
Unemployed	0.04 (0.19)	0.12 (0.32)	0.16 (0.37)
In current home for 5+ years	0.82 (0.38)	0.34 (0.48)	0.64 (0.48)
Moved since last year	0.05 (0.22)	0.25 (0.43)	0.1 (0.3)
Weighted N	17,107,944	7,585,030	4,390,937

Statistics are computed from the 2018 full Census microdata. The sample is restricted to the reference individual for all households in each of the three tenure modes, excluding respondents in non-traditional or institutional housing. Estimates use sampling weights provided by the national statistical institute INSEE, and provide the mean and (in brackets) standard deviation for each variable.

ence individual for households in three subgroups (owner-occupiers, renters in the private sector, and renters in the public sector), obtained from exhaustive Census data. Private sector renters live in smaller households, are younger, more likely to be full-time employed and less likely to be retired, and are overwhelmingly more likely than either owners or public sector renters to have moved in recent years and to have lived in their home for a short time period.

All of these observable characteristics are also associated with a more substantial propensity to move across municipalities. Nonetheless, when estimating a linear probability model for having moved from a different municipality in the last year on a variety of individual characteristics, as presented in table 1.E.2, the coefficient on renting remains high, even after controlling for a number of individual covariates. The partition of the population of households alongside characteristics associated with spatial mobility indicates a clear segmentation between the population of mobile private sector renters and mostly immobile owner-occupiers.

Table 1.E.2: Probability of having moved across cities

	(No controls)	(With controls)
Owner-occupier		(Reference category)
Private sector renter	0.14****	0.07****
	(0.00)	(0.00)
Public sector renter	0.01****	-0.01****
	(0.00)	(0.00)
R-Square	0.05	0.09
Observations (unw.)	8762383	8762383

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Regressions coefficients are obtained from a linear probability model using the 2018 full Census microdata. Controls in the second column include household size, an indicator for being full-time employed, an indicator for being out of the labor force or retired, and a quadratic in the age of the reference individual. The sample is restricted to the reference individual for all households in each of the three tenure modes, excluding respondents in non-traditional or institutional housing, corresponding to a weighted number of observations of $N^w = 29,083,911$. Estimates use sampling weights provided by the national statistical institute INSEE, and robust standard errors to heteroskedasticity.

Spatial variation in the propensity to move is closely associated with the presence of private rental housing in a given location, as documented above for the cross-section of French municipalities. This fact also holds over time within cities. Using panel data from 2009 to 2018 at the municipality-year level obtained from the French Housing survey matched to Census data, I show in table 1.E.3 that even conditioning on municipality fixed-effects, a higher share of private rental housing is tightly associated with increases in measures of spatial and housing mobility. This pattern is not claiming a direct causal link from the availability of rental housing to the evolution of household mobility, but the existence of a residual correlation within municipalities over time provides suggestive evidence that the close connection between the two phenomena captures more than time-invariant structural characteristics of a locality.

Table 1.E.3: Recent moves and private rental share

	Pooled	Year FE	City FE	2-way FE
Share of private rentals	0.363****	0.363****	0.150****	0.162****
	(0.0327)	(0.0328)	(0.00997)	(0.0107)
R-Square	0.572	0.573	0.936	0.937
Observations	359390	359390	359388	359388
Clusters	36682	36682	36680	36680
Municipality FE	No	No	Yes	Yes
Year FE	No	Yes	No	Yes

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Regressions coefficients are shown for the association between the share of private rental housing in a municipality and the share of households who moved in less than two years ago. They are obtained from a panel regression with city and year fixed-effects, across France's municipalities, using data from the French Housing survey matched to the Census for 2009-2018. Estimates use robust standard errors clustered at the municipality level, and are weighted by city population.

1.E.2 Housing investor characteristics

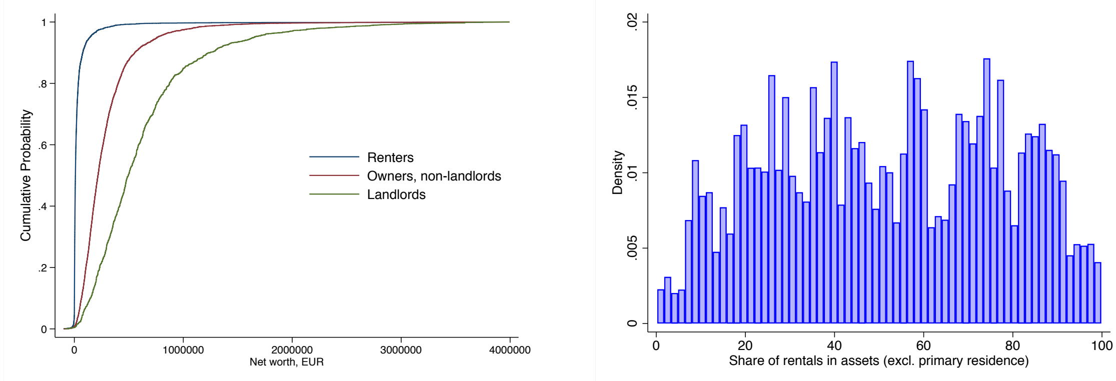

(a) Distribution of net worth by landlord status

(b) Share of rental units in gross wealth

Figure 1.E.2: **Landlords wealth and limited diversification**

Panel a plots the cumulative distribution function of net worth for households in three groups: renters, owner-occupiers who do not own any rental dwelling, and landlords defined as households who own at least one unit that they rent out for the entire year. Panel b plots the distribution of landlords according to the portfolio share of rental units in their wealth, excluding the value of their primary residence, using matched household-assets data from the wealth survey. The median portfolio share of rental units in gross wealth is 52 percent. Data are computed at the household level, from the *Enquete Patrimoine* 2017-2018, a wealth and income survey matched to summary information from household tax returns which oversamples high wealth and high income individuals.

To motivate the segmentation between the saving behavior of landlords and the rest of the population, I show using microdata from a representative survey of households assets that landlords tend to differ substantially from both other owner-occupiers, and from renters, along numerous characteristics. Using wealth survey data from 2018, table 1.E.4 provides some demographic information on the reference individual for households in three subgroups (landlords, non-landlord owner-occupiers, and renters), obtained from the *Enquete Patrimoine*. Landlords have higher gross wealth, receive higher disposable income, are older and more frequently married, and tend to live in larger households than either renters or non landlord owner-occupiers. Both types of owners have on average lived for much longer than renters in their current residence location. Figure 1.E.3 shows that the income distribution of landlords tend to first-order stochastically dominate the income distribution of both non-landlord owners and of renters.

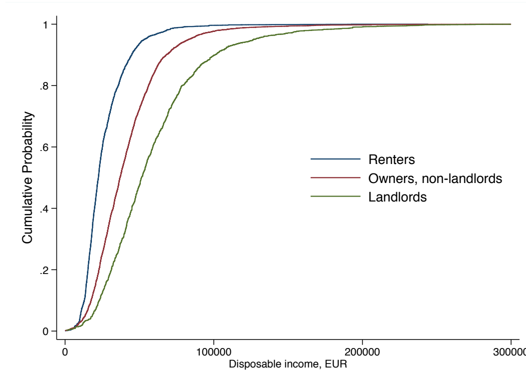

Figure 1.E.3: **Distribution of income by landlord status**

The figure documents the higher disposable income of landlords relative to other owner-occupiers who do not own rental housing and to renters. Data are computed at the household level, from the *Enquete Patrimoine* 2017-2018, a wealth and income survey matched to summary information from household tax returns which oversamples high wealth and high income individuals.

Table 1.E.4: Characteristics of households by landlord status

	Landlords	**Non-LL owners**	**Non-LL renters**
Household size	2.52	2.38	2.05
	(1.26)	(1.24)	(1.33)
Age of reference individual	55.8	57.6	48.8
	(14.6)	(16.4)	(18.1)
Gross wealth	821,517	352,137	41,878
	(1,605,048)	(897,508)	(172,445)
Disposable income	60,705	41,418	26,252
	(53,649)	(30,443)	(15,477)
Married	0.74	0.67	0.34
	(0.44)	(0.47)	(0.47)
Owns more than one dwelling	1	0.11	0.05
	(0)	(0.31)	(0.23)
Years in current dwelling (med.)	15	17	6
	(14.9)	(15.8)	(11.0)
Weighted N	2,958,885	13,988,108	11,209,854

Statistics are computed from the 2018 *Enquete patrimoine* wealth survey. The sample is restricted to the reference individual for all households in each of the three subgroups, excluding respondents in non-traditional or institutional housing. Estimates use sampling weights provided by the wealth survey, and provide the mean and (in brackets) standard deviation for each variable.

1.F Additional results

1.F.1 Retiming and bunching evidence

I evidence graphically additional margins of adjustment to conditional housing subsidies. First, a high-frequency analysis of the timing of both developer requests of building licenses, and signing dates of actual transactions, demonstrates the salience and key role of tax motives in triggering housing supply decisions. Second, households bunch the purchase price of their new housing acquisitions below a kink in the subsidy schedule, providing evidence of sensitivity to changes in marginal subsidy rates.

End-of-year abnormal investment behavior: sales Abnormal end-of-year bunching in transactions is observed in the last month of each fiscal year for "off-plan" sales (denoted as *VEFA*, which correspond to pre-sales of units still under construction or for which construction has not started yet), even in the absence of any change in eligibility or tax treatment, as shown in figure 1.F.1.

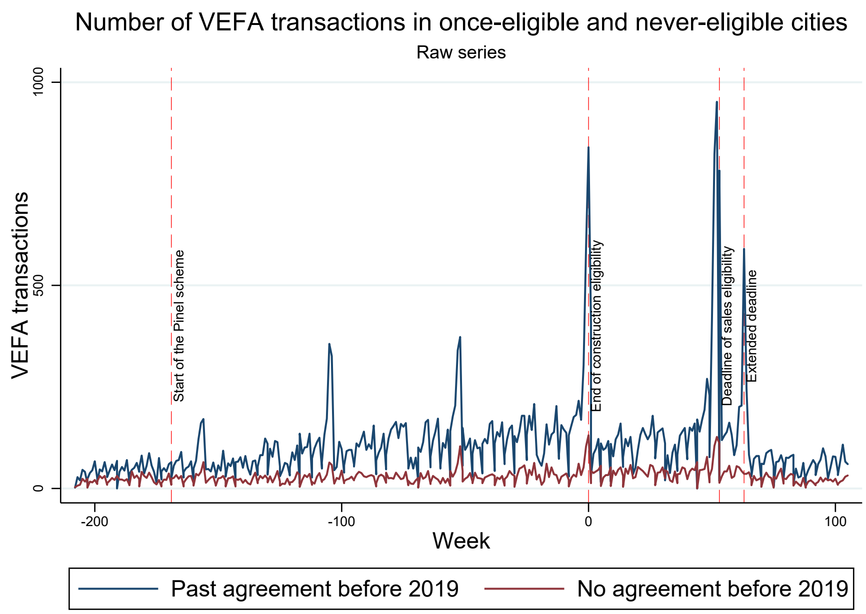

Figure 1.F.1: **Total number of "off-plan" sales in B2+ and B2- municipalities**

The figure plots the weekly number of sales identified with an "off-plan" dummy in France from 2014 to June 2020, respectively in never-eligible B2- (red line) and once-eligible B2+ (blue line) municipalities. Sources: *DVF* database, an exhaustive record of housing transactions in France made available by the French Finance Ministry. Sales restricted to houses and flats with an available transaction value - see appendix 1.C for a detailed description of the data.

This provides evidence of the salience of the tax subsidy, but also potentially of a rational time-shifting of individual investment decisions in order to minimize the time elapsed between the investment in residential housing, and the claiming of the tax credit rebate during the next fiscal year. This behavior is particularly pronounced in anticipation of changes to the value of the subsidy in B2+ locations, and represents the household analog of the end-of-year bunching of tax-minimizing investments documented by Xu and Zwick (forthcoming) for the case of US firms subject to the corporate income tax. Sales in eligible B2+ cities exhibit sharp monthly spikes at the initial planned eligibility limit (December 2017), and at the interim (December 2018) and final (March 2019) eligibility deadlines, providing further credible evidence of the causal impact of the subsidy on new home sales.[81] These spikes are statistically and economically significant: in the last week of 2018, new home sales in the

[81] Consistent with the findings of this paper, Berger, Turner, and Zwick (2020) and Best and Kleven (2018), studying temporary home-ownership subsidies in the United States and transaction taxes exemptions in the United Kingdom, respectively, also find significant evidence of time-bunching right before eligibility to each of the policies ended.

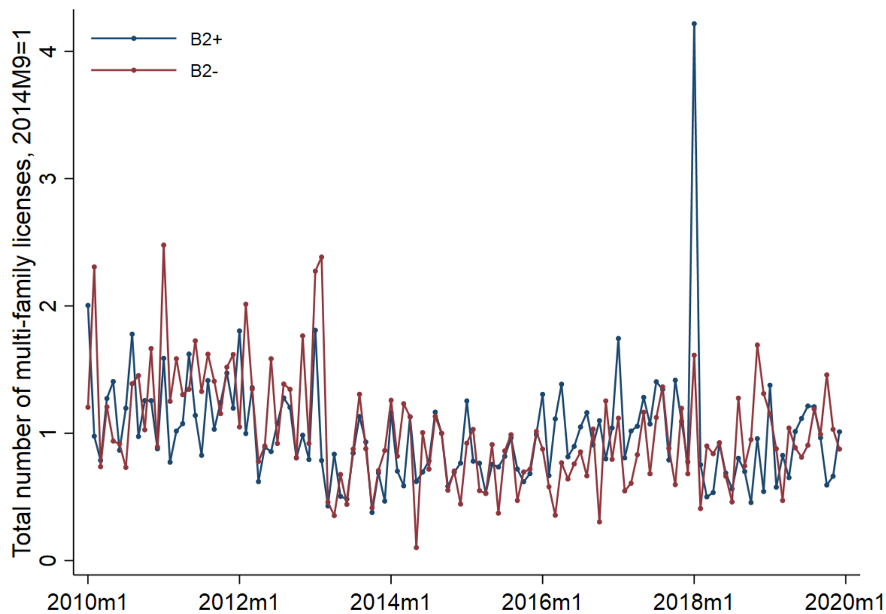

Figure 1.F.2: **Total number of building license requests in B2+ and B2- municipalities**

The figure plots the monthly number of building license requests for multi-family units, respectively in never-eligible B2- (red line) and once-eligible B2+ (blue line) municipalities. Sources: *Sit@del* database, an exhaustive record of building license requests - see appendix 1.C for a detailed description of the data.

900 eligible B2 municipalities reached more than EUR 200 million, roughly twenty times the average weekly new home sales in both eligible and ineligible cities after the policy ended.

End-of-eligibility re-timing: licenses For housing building permits granted to developers, the reform of the Pinel scheme (which was passed in Parliament in November 2017) allowed for housing sales in B2+ cities to remain eligible only if a building license for the unit had been *requested* before December 31, 2017. In response to this incentive, the number of building licenses requested by developers in eligible B2+ cities spiked in December 2017. I evidence in figure 1.F.2 the presence of this spike in developer requests of building licenses immediately before the planned end of eligibility, using detailed information from the Sit@del2 (building permits requests) database.

Kinks in the subsidy schedule: purchase price Finally, exploiting exhaustive individual tax returns information, I evidence the presence of bunching at the kink in the subsidy schedule arising from the phase-out of the subsidy for units above 300,000 €. Figure 1.F.3 plots the

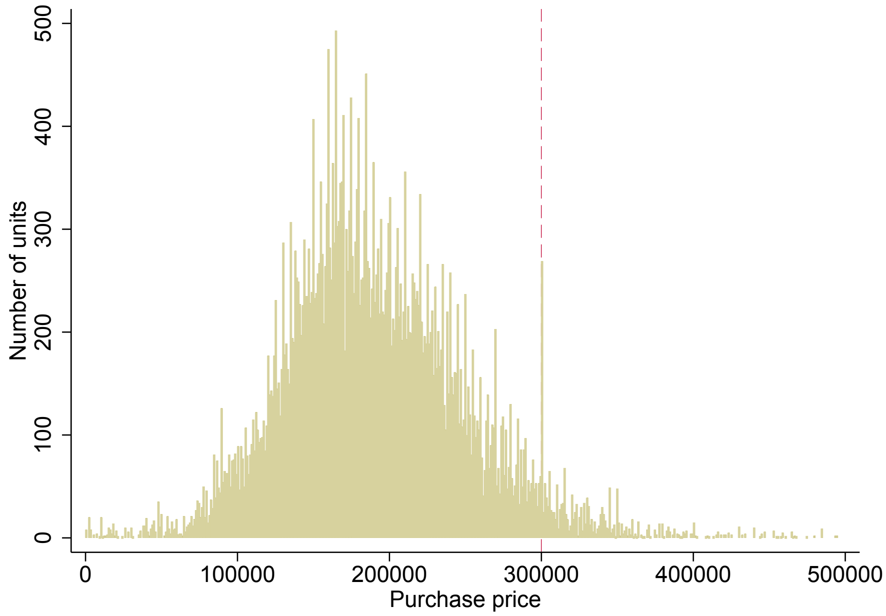

Figure 1.F.3: **Purchase price of units under the Pinel scheme**

The figure plots the distribution of purchase prices for units purchased under the Pinel scheme in the year 2016. Sources: *POTE* database, an exhaustive record of individual tax returns - see appendix 1.C for a detailed description of the data.

distribution of acquisition prices for all units under the *Pinel* tax schedule and evidences a clear spike exactly at 300,000 €, the value at which the proportional tax credit is maximized and above which the marginal subsidy is reduced to 0.

1.F.2 Additional local effects of eligibility

Land sales A common concern surrounding housing subsidies is the capitalization of incentives in inelastically supplied local input prices, most notably land (Carozzi, Hilber, and Yu, 2019; Bono and Trannoy, 2019). Such input price rises push up construction costs as a consequence of housing subsidies. Using the DV3F database, I restrict the sample to sales of constructible land (so-called "*terrains a batir*").

Figure 1.F.4 plots the total monthly value of land sales in eligible B2 municipalities and their counterparts throughout the policy period. After following parallel trends to ineligible municipalities for much of 2014, the value of land sales in B2 towns with an agreement rose sharply throughout the policy eligibility period, especially in 2017, before the spike in

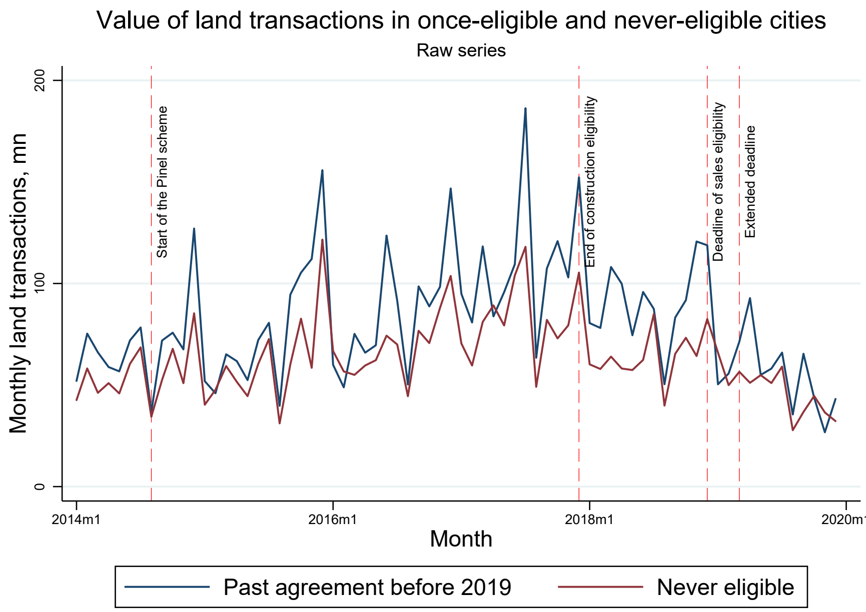

Figure 1.F.4: **Total value of constructible land sales in B2 municipalities**

The figure plots the monthly value of constructible land sales (million euros) in France from 2014 to June 2020, respectively in never-eligible (red line) and once-eligible (blue line) B2 municipalities. Sources: *DVF* database, an exhaustive record of housing transactions in France made available by the French Finance Ministry. Sales restricted to houses and flats with an available transaction value - see appendix 1.C for a detailed description of the data.

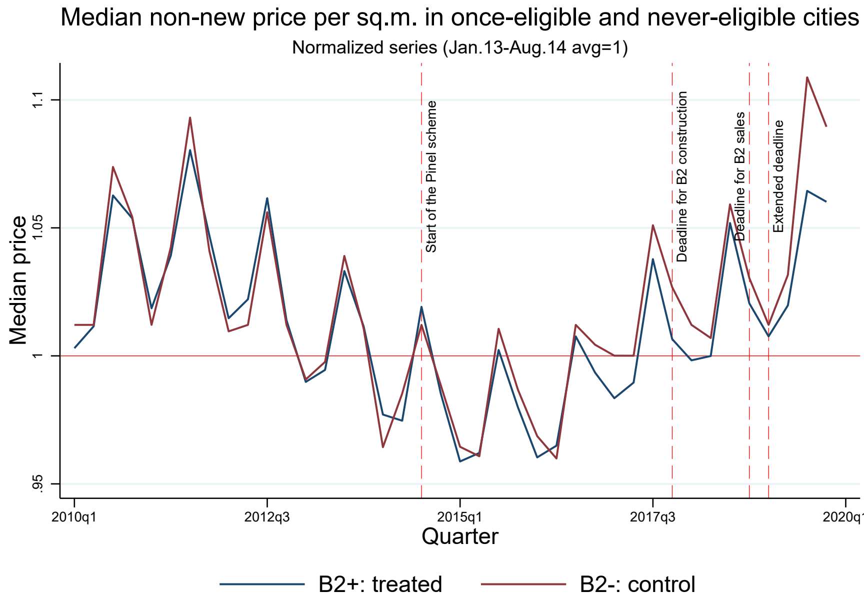

Figure 1.F.5: **Median price of existing units in B2 municipalities**

The figure plots the median price per square meter in transactions for *existing* units, respectively in never-eligible (red line) and once-eligible (blue line) B2 municipalities. Sources: *DVF* database, an exhaustive record of housing transactions in France made available by the French Finance Ministry. Sales restricted to houses and flats with an available transaction value - see appendix 1.C for a detailed description of the data.

licenses granted. After the end of the policy, the value of land sales in formerly subsidized B2 towns falls back to a trend and level comparable to ineligible cities.

Prices of existing units As stated by Kotlikoff (1983), a specific feature of investment tax incentives is their effect on the price of "old" capital: "*since equally productive units of new and old capital must sell for the same price, tax provisions favoring new capital imply a lower price for existing capital*". While the only existing test of this hypothesis, to my knowledge, focuses on the price for used capital goods purchased by corporations (Edgerton, 2011), in my setting, observing local transactions for existing units in treated and untreated locations allows me to test for the presence of price effects for existing dwellings, i.e. location-specific used residential capital. Figure 1.F.5 displays the normalized median price per square meter of existing units in B2+ and B2- municipalities over time, and provides suggestive evidence that treated locations experience a relative decline in the price of existing homes, as competing substitutable capital reduces their expected rental value.

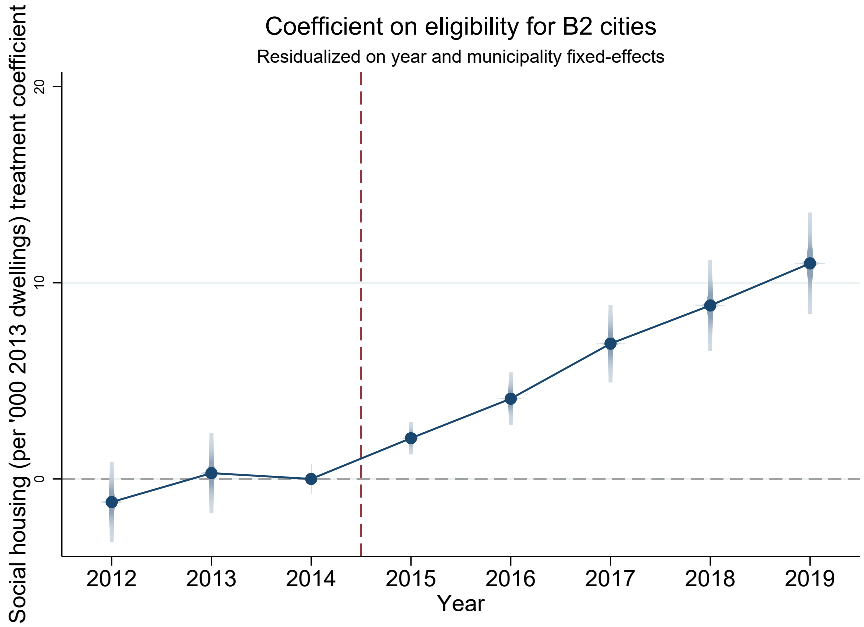

Figure 1.F.6: **Effect on social housing units**

The figure plots the estimated effect of eligibility to the *Pinel* scheme on social housing units. It plots, for B2+ (treated) relative B2- (control) towns, the coefficient on eligibility in equation 1.8 for the number of housing units per thousand of 2013 dwellings in the town. Standard errors are clustered at the municipality level and 95% confidence intervals are shown. Sources: *RPLS* database, an exhaustive record of social housing units in France made available by the French Housing Ministry – see appendix 1.C for a detailed description of the data.

Social housing units Using administrative data on the universe of social housing units provided by the French government, I also quantify a positive induced effect on social housing. Figure 1.F.6 displays the results of the event-study regression coefficient on eligibility around entry. Part of the mechanism for the observed increase in social housing units stems from economies of scale, through which some of the individual units in a multi-family building targeted by developers for buy-to-let individual investors are purchased by local governments and public-private partnerships (*bailleurs sociaux*) to provide public housing in mixed-income communities.

1.G Additional appendix figures

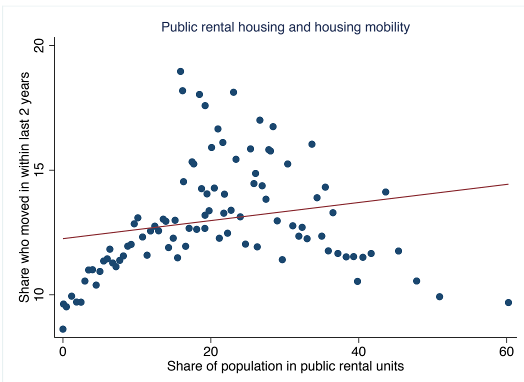

(a) Social rental housing and housing turnover

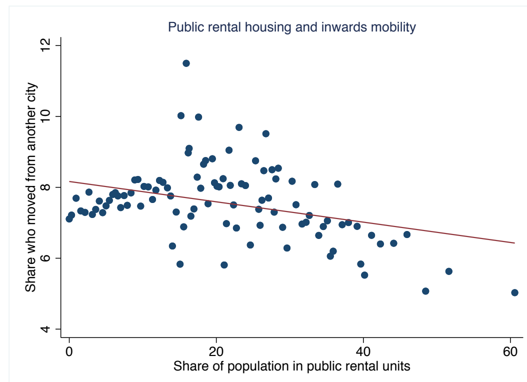

(b) Social rental housing and inwards mobility

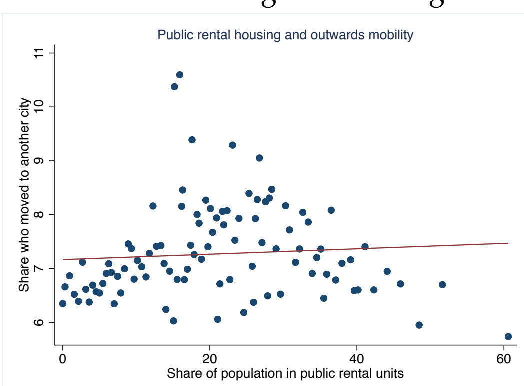

(c) Social rental housing and outwards mobility

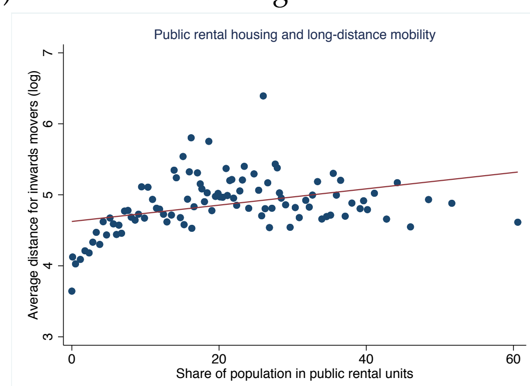

(d) Long-distance moves and social rentals

Figure 1.G.1: **Public rental housing and spatial mobility**

Panel a plots the (population-weighted) relationship between the share of the population living in public rental units in a municipality, and the share of the population living in units which they moved in less than two years earlier, in 2018. Panel b (resp c) plots the (population-weighted) relationship between the share of the population living in public rental units in a municipality, and the share of the population which moved into (resp. out of) the municipality in the last year, in 2018. Panel 1.G.1d plots the (flow-size weighted) average distance of moves ($\bar{d_i} = \sum_j \frac{F_{ij}}{\sum_j F_{ij}} d_{ij}$) by inwards movers from j to to municipality i in the last year (in logs), against the share s_i of the population living in public rental units in i, in 2018. The distance of moves is computed as the haversine distance between the centroids of the municipality of origin and the destination municipality. All panels present binned scatter plots where the average value of the y-axis variable is plotted against each (population-weighted) centile of the x-axis variable, and data for these panels come from France's national statistical institute *Base Logements*, a housing survey matched to the Census and combined with register data.

(a) Rent-price ratios and prices: across municipalities

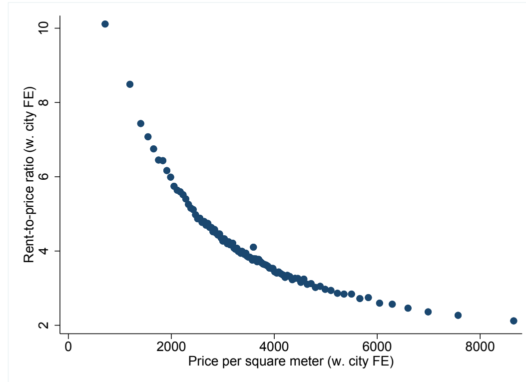

(b) Rent-price ratios and prices: within municipalities

Figure 1.G.2: **Rent-to-price ratios against prices**

Panel a plots a binned scatter plot of rent-to-price ratios against median purchase price per square meter across municipalities, where y-axis values are averaged by centile bins of the x-axis variable. The rent-price ratios are computed using a cross-section of rents per square meter for multi-family units in 2018 made available by the French Housing secretariat at the municipality level, and median purchase prices for multi-family units from 2017 to 2019 using exhaustive transaction deeds data from the *DV3F* database. They only include municipalities with more than 11 transactions for multi-family units over the period. Panel b plots a binned scatter plot of rent-to-price ratios against purchase price per square meteracross rental units purchased under the *Pinel* subsidy scheme, for municipalities with more than one unit, where both variables are demeaned by municipality, and y-axis values are averaged by centile bins of the x-axis variable. The rent-price ratios are computed using individual data on all c. 65,000 purchases made in 2016 under the *Pinel* scheme, exploiting buyers' commitment forms (2044-EB schedules attached to individual tax returns) which record the initial value of the rent under the affordable lease, as well as the purchase price and floor area of the unit to compute the taxpayers' tax reduction.

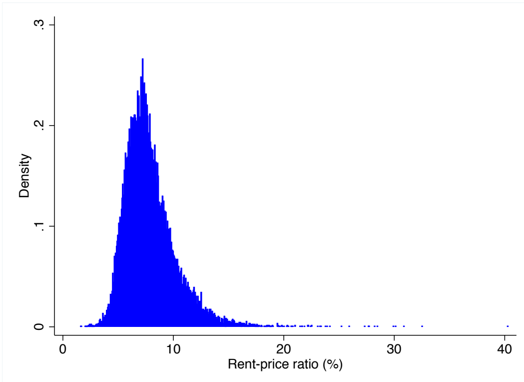

(a) Rent-price ratios: all municipalities

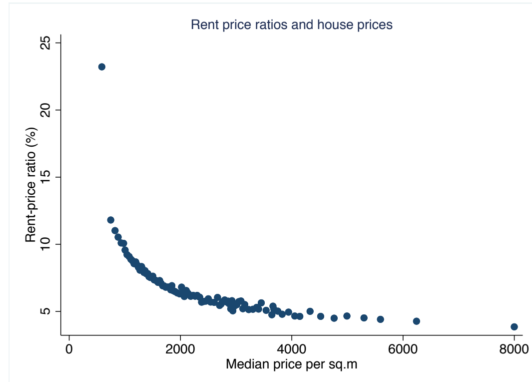

(b) Rent-price ratios and prices: all municipalities

Figure 1.G.3: **Rent-to-price ratios across all municipalities**

Figure 1.G.3 documents the cross-city dispersion in rent-to-price ratios. It is constructed similarly to figure 1-2 but includes all municipalities with available data and does not restrict to municipalities with more than 11 transactions. Panels 1.G.3a plots the distribution of rent-to-price ratios across municipalities. Panel 1.G.3b plots a binned scatter plot of rent-to-price ratios against median purchase price per square meter across municipalities. The figures are computed using a cross-section of rents per square meter for multi-family units in 2018 made available by the French government at the municipality level, and median purchase prices for units from 2017 to 2019 using exhaustive transaction deeds data from the *DV3F* database.

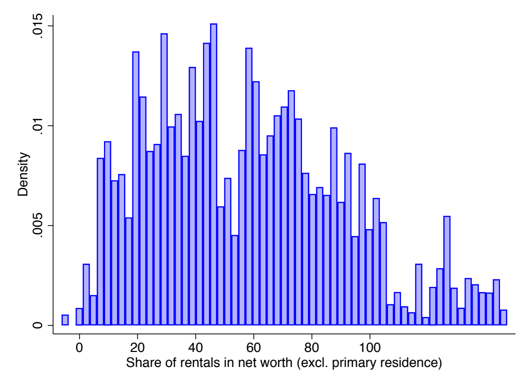

Figure 1.G.4: **Portfolio share of rental units in net worth (excl. primary residence)**

Figure 1.G.4 documents the undiversified nature of housing investment in rental units. It plots the distribution of landlords according to the portfolio share of rental units in their net wealth, excluding the value of their primary residence, using matched household-assets data from the wealth survey. The median portfolio share of rental units in gross wealth is 52 percent. Data are computed at the household level, from the *Enquete Patrimoine* 2017-2018, a wealth and income survey matched to summary information from household tax returns which oversamples high wealth and high income individuals.

(a) New home sales

(b) New homes listed

Figure 1.G.5: **New home sales and new listings**

Panel (a) plots the quarterly number of sales of newly built homes (normalized to 1 in 2014Q3) in B1 (blue line), B2 (red line), and C (green line) municipalities. Panel (b) plots the quarterly number of new listing of newly built homes (normalized to 1 in 2014Q3) in B1 (blue line), B2 (red line), and C (green line) municipalities. Sources: *ECLN* database, an exhaustive record of new home sales transactions in France made available by the French Housing Ministry - see appendix 1.C for a detailed description of the data.

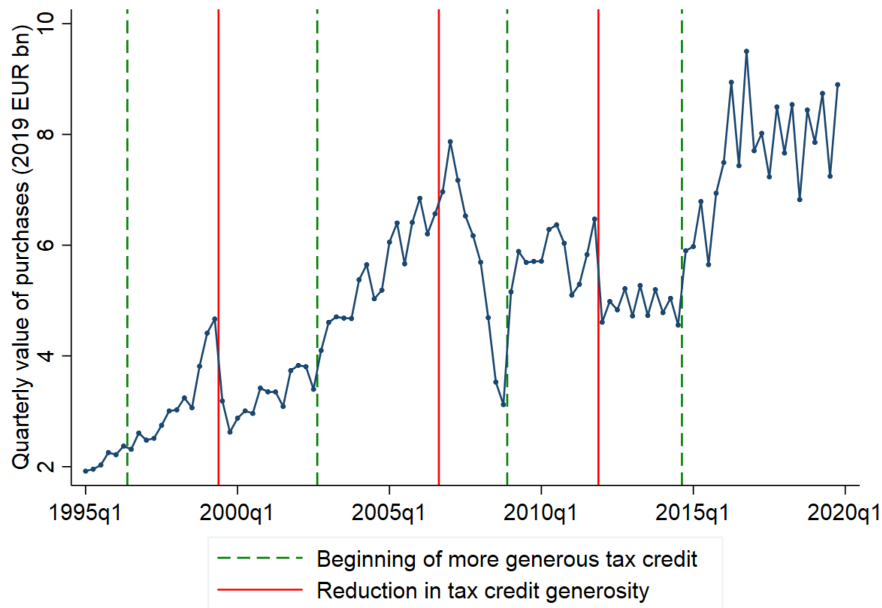

Figure 1.G.6: **The tax responsiveness of residential investment**

The figure plots the total quarterly value (in 2019 billions of EUR) of new housing units purchased from developers in France from 1995 to 2019. Dashed vertical green lines indicate the beginning of a more generous investment tax credit or bonus depreciation scheme for new housing; continuous vertical red lines indicate a switch to a less generous tax regime. The successive tax regimes are described in table 1.B.4. The total value of units sold is computed from microdata in the *ECLN* database, an exhaustive survey of developer-led projects of five or more units.

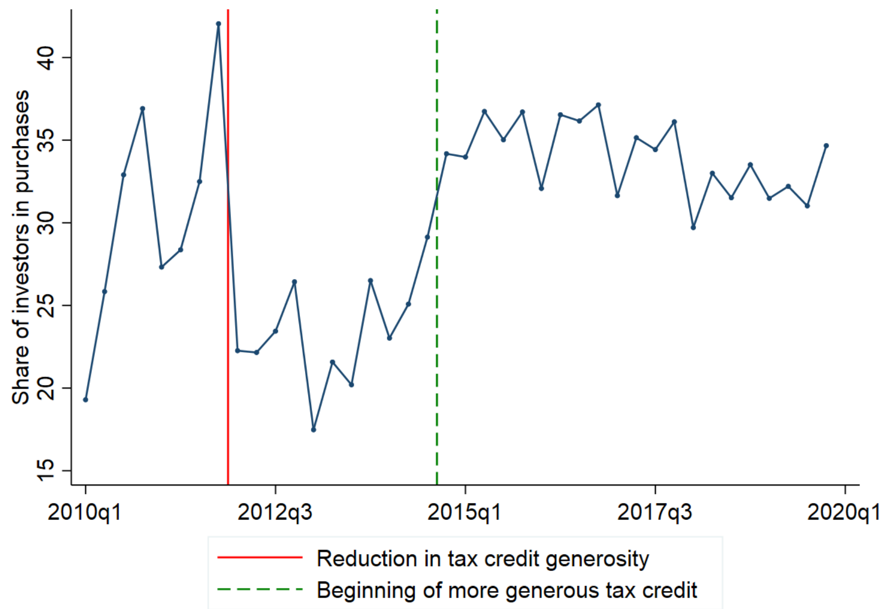

Figure 1.G.7: **The tax responsiveness of buy-to-let investment**

The figure plots the share of investors (as opposed to owner-occupiers) in the value of new housing units purchased from developers in France from 2010 to 2019. The continuous vertical red line indicates the end of the more generous version of the *Scellier* scheme in 2011, and the switch to a lower level of the tax credit. The dashed vertical green line indicates the beginning of the more generous *Pinel* investment tax credit in October 2014. The successive tax regimes are described in table 1.B.4. The share of units sold to investors is computed from microdata in the *ECLN* database, an exhaustive survey of developer-led projects of five or more units.

Chapter 2

Specializing in Density [1]

2.1 Introduction

Does the distribution of economic activity *within* countries affect the pattern of trade *across* countries? Most analyses of comparative advantage treat countries as unified factor markets or equilibrium "points" in the production space. However, there is mounting evidence that the *distribution* of of factors within countries—and in particular, population density—is a key determinant of productivity.[2] The impacts of place-based policies, of urban agglomeration, and of spatial sorting, as well as their implications for domestic welfare and inequality, are the subject of substantial analysis.[3] Urban planners and politicians debate the role of density in shaping features of life ranging from firm location decisions to local pollution to violent crime. This chapter documents that domestic heterogeneity also has a major impact on patterns of trade—differences across countries in the extent of agglomeration (or lack thereof) shape comparative advantage.

The hypothesis domestic heterogeneity could affect comparative advantage is general, and a version of it dates back to Courant and Deardorff (1992), who argue that the "lumpiness" of factor distribution can affect a country's pattern of exports. We take it directly to

[1]This chapter is joint work with Jacob Moscona. We thank Gilles Duranton, Teresa Fort, Guy Michaels, Nathan Nunn, Ben Olken for advice and comments. We are grateful to seminar participants at MIT, Dartmouth, and the Urban Economics Association for helpful feedback.

[2]For example, see Keesing and Sherk (1971), Ciccone and Hall (1996), Duranton and Puga (2004), and Moretti (2012) and more recently Davis and Dingel (2014) and Gaubert (2018).

[3]See e.g. Hsieh and Moretti (2019) and Kline and Moretti (2014b)

data by investigating one particular but central example of domestic heterogeneity: population density and distribution. Density may boost productivity through several potential mechanisms; dense cities ease search and matching frictions in the labor and product market, attract high-skilled and talented workers, provide large and local consumption markets, and serve as hubs for high-tech investment and innovation (e.g. Duranton and Puga, 2004; Moretti, 2012). Crucially, moreover, and as we will document in detail, density bolsters productivity differentially across industries, some of which end up located at the center of large agglomerations while others end up in smaller cities or sparsely populated areas.[4]

This logic suggests that the extent to which a country's population is concentrated in dense areas might affect not only its domestic productivity, but also its international specialization. If industries benefit differently from population density, holding all other country-level characteristics constant, countries with a more concentrated population distribution will have a comparative advantage in industries that benefit disproportionately from agglomeration. While a range of work has analyzed the effect of trade on domestic economic geography, this reverse relationship—the effect of patterns of urbanization on patterns of trade—has received little attention.[5] Using novel measures of industry-level "density affinity" and country-level "population concentration," we show that urban density is a significant determinant of international exports.

We first present a model that illustrates how the distribution of factors of production within countries—i.e. having a concentrated versus dispersed population—affects patterns of trade. In the model, industries vary in the extent to which they benefit from the population density of the location in which production takes place.[6] Countries are composed of locations endowed with different sector-neutral productivities. Endogenously, countries with more dispersion in sector-neutral productivity exhibit higher population-weighted density (i.e. a more concentrated population) and have a comparative advantage in sectors that benefit relatively more local agglomeration.

[4]On the impact of density on *sector-specific* productivity and role of density in determining heterogeneity across sectors in spatial sorting, see also Nakamura, 1985; Rosenthal and Strange, 2004; Faggio, Silva, and Strange, 2017.

[5]On the impact of trade shocks on domestic economic geography, see, for example: Autor, Dorn, and Hanson (2013), Caliendo, Dvorkin, and Parro (2015), Dix-Carneiro and Kovak (2015), Ramondo, Rodríguez-Clare, and Saborío-Rodríguez (2016), and Bakker (2018)

[6]In the baseline model, we are agnostic about the source of this variation in agglomeration externalities.

The theory provides three key insights. First, motivated by evidence that countries display significant domestic spatial heterogeneity in factor prices, product specialization, and relative productivity (e.g. Porter, 2003; Desmet and Rossi-Hansberg, 2013), our model formalizes the idea that the relevant units of observation for understanding comparative advantage are regions within countries where production takes place. This is different from most models of comparative advantage, which focus on aggregate country-level characteristics that are taken as given. Second, our model documents how regional variation can be aggregated to uncover country-level determinants of comparative advantage. The model formalizes our use of "population-weighted density" as the country-level summary of within-country heterogeneity in population density. Finally, the model provides theoretical justification for our main empirical framework and result: countries with higher population-weighted density have relatively lower autarky prices in sectors that benefit relatively more from agglomeration; hence, their exports exhibit a revealed comparative advantage in these sectors.

We then empirically investigate whether the distribution of population within countries is an important determinant of comparative advantage. Our empirical strategy requires two main ingredients: (i) a sector-level estimate of "density affinity," or the extent to which production in each sector is disproportionately located in denser locations, and (ii) a country-level estimate of population concentration.

To measure industry-level density affinity, we turn to detailed business location data across US urban areas from the County Business Patterns (CBP) and non-parametrically estimate the extent to which each sector is disproportionately located in denser locations. To account for potential endogeneity in the correlation between density and industry specialization, we use subterranean geological instruments that exogenously shift local density independently from other city-level characteristics by easing vertical construction costs and constraints. This generates causal estimates of the marginal impact of population density on industry-level production. In the end, this procedure yields industry-level measures of density affinity across all 4-digit NAICS manufacturing sectors; the substantial heterogeneity in density affinity that we estimate lends credibility to the modeling assumption of significant variation in sector-specific sorting with respect to population density.

To measure population-weighted density across regions and countries, we rely on satellite-derived gridded population data from the *LandScan* database. *LandScan* incorporates comprehensive country-level census data on the distribution of population, and derives gridded population estimates using "smart interpolation," a multi-layered, asymmetric, spatial modeling approach. These data make it possible to estimate characteristics of the geographic population *distribution* of each country. To measure population-weighted density, we sum population density across grid cells within each country, weighting each cell by its total population. This captures the experienced population density of the average person in the country and measures the concentration of population across space.

Before turning to cross-country trade, we focus on the exporting patterns of US States. Using the *LandScan* data, we estimate the population-weighted density of each state, and document that denser states indeed export relatively more in "density-loving" sectors.[7] While this result is a preliminary test of our hypothesis, it also validates our density affinity measures as supply side determinants of sector productivity, rather than the product of path dependence or demand-side forces. Our estimates of density affinity from US data could have been driven by the fact that certain sectors are over-represented in certain US locations for historical or demand-side reasons; however, the state-level export results suggests that density-loving sectors are indeed more productive in denser regions within the US. Population concentration is a source of state-level comparative advantage.

Next, we investigate the role of density as a source of country-level comparative advantage. Country-level estimates of population weighted density are displayed in the map in Figure 2.1.1. There is substantial variation in density across countries, even within continents and income levels. For example, Finland and Sweden are two of the wealthiest and also two of the least dense countries in the world, by our measure; indeed, both countries have strong revealed comparative advantage in pulp and paper product exports, one of the least density-loving sectors.[8] Within sub-Saharan Africa, Botswana is among the least dense countries while the nearby Democratic Republic of Congo and Djibouti, among the world's poorest countries, are among the densest.[9] Finally, the United States has mid-range

[7] While some recent studies have attempted to estimate export data at the metropolitan level (see e.g. the database constructed by Tomer and Kane, 2014), most trade flows data are still collected at a broader level of aggregation. The lowest level of consistent and exhaustive trade reporting in the United States is the state.

[8] See Sweden exports in the Atlas of Economic Complexity for HS4 codes 4800-4810, NAICS code: 3221.

[9] Indeed, Djibouti, exhibits a strong revealed comparative advantage in semiconductors, one of the most density-loving sectors. See Djibouti exports in the Atlas of Economic Complexity for HS4 code: 8541, NAICS

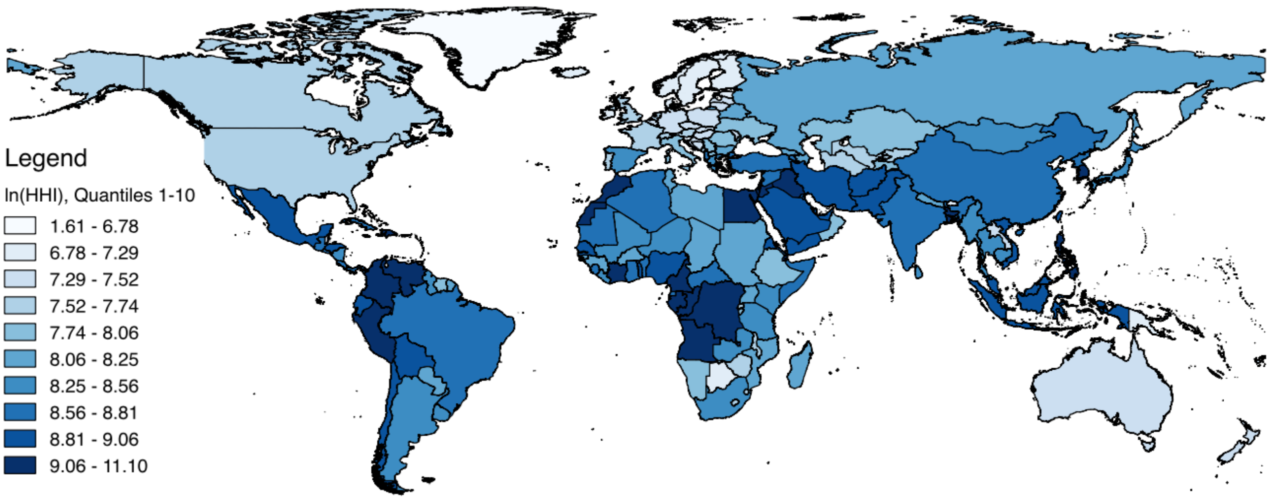

Figure 2.1.1: **Population weighted population density across countries (deciles)**

The figure is a map in which countries are color-coded based on their population-weighted density decile. Darker countries have higher population-weighted density.

population-weighted density since it has both very dense cities, as well as a relatively large share of the population living in suburbs, towns, and rural areas.

We systematically investigate the relationship between population density and comparative advantage and show that countries with higher population-weighted density have a revealed comparative advantage in density-loving sectors. This finding is robust to the inclusion of a broad range of country and industry-level controls, including the skill and capital intensity of each sector, as well as country-level income, skill endowment, specialization in agriculture, and other covariates that might bias the relationship between density and exports. The results are also similar across a range of possible parameterizations of the density affinity measure; either including or excluding observations with zero trade; and using either OLS or Poisson pseudo-maximum likelihood estimators.

To correct for potential reverse causality from trade flows to density (see Krugman and Elizondo, 1996; Ades and Glaeser, 1995), we exploit differences in states' and countries' *historical* population and city size distributions to construct instruments for modern density.

code: 3344.

Data on the global distribution of cities and their populations for historical periods were collected by Chandler (1987), and recently digitized by Reba, Reitsma, and Seto (2016). While patterns of trade might affect modern economic geography, it is unlikely that modern patterns of trade, which have evolved substantially in recent decades and particularly after the Second World War, affected the *historical* (c. 1900) distribution of cities within countries (see Irwin, 2017). Using this strategy, the estimated effect of density on trade flows from our baseline results remains very similar. In our sample of countries, we find that the impact of the within-country population distribution on patterns of trade is comparable to and if anything slightly larger in magnitude than the impact of human or physical capital.

Finally, we investigate potential channels underpinning the relationship density affinity and trade. Our goal was to capture all possible effects of population density on industry-level productivity; therefore, our baseline density affinity measure does not take a stand on any particular mechanism. However, using data on the task content of production in each industry, we find that the relative importance of different tasks in more vs. less "density loving" sectors is an important driver of our findings; in particular, denser countries tend to have a comparative advantage in sectors that rely on more interactive and collaborative tasks, while less dense countries specialize in sectors that rely on interaction with machines. We also find evidence that the research and development (R&D) intensity and natural resource input share of each industry are additional intervening mechanisms, consistent with evidence that dense cities facilitate and spur innovation (e.g. Duranton and Puga, 2001; Duranton and Puga, 2004; Moretti, 2012) and that only industries that do not rely on natural resources are free to locate in cities (Ades and Glaeser, 1995). Last, we find no evidence that the results are driven by industry-level skill or capital intensity, or reliance on service-sector inputs (see e.g., Abdel-Rahman, 1994). Together, the mechanisms that we propose explain about 65% of our baseline estimate, suggesting that additional and un-observed industry characteristics also contribute to industries' sorting and resulting international specialization.

This study is at the intersection of several areas of research. Our theory is most closely related to Courant and Deardorff (1992) and Courant and Deardorff (1993), who argue that patterns of trade come not only from relative factor abundance, but also from factor distribution ("lumpiness"). The idea has been explored more recently by Debaere (2004), Bernard, Robertson, and Schott (2010), and Brakman and Van Marrewijk (2013). We directly measure the "lumpiness" of population density from satellite data and estimate the impact of population distribution on patterns of trade.

This chapter also builds on prior work studying the sorting of sectors across cities (e.g. Davis and Dingel, 2014; Gaubert, 2018) and the differential extent of agglomeration across sectors (Nakamura, 1985; Rosenthal and Strange, 2001; Rosenthal and Strange, 2004; Holmes and Stevens, 2004; Ellison, Glaeser, and Kerr, 2010; Faggio, Silva, and Strange, 2017). We extend work in this area by developing a new strategy to estimate industry-specific sorting with respect to density and investigate the relationship between this within-country sorting and cross-country trade. This chapter also builds on work devoted to understanding the causes and consequences of city size distributions and Zipf's law (e.g., Gabaix, 1999a; Gabaix, 1999b; Ioannides and Overman, 2003). Using gridded data, we observe the full population distribution—both inside and outside administratively defined urban areas—-and measure the average experienced density of each country, which differs markedly around the world, even conditional on total population and average density or urbanization

A broad set of work studies the interplay between trade and within-country heterogeneity, often by by highlighting the effect of international trade on within-country disparities (Autor, Dorn, and Hanson, 2013; Caliendo, Dvorkin, and Parro, 2015; Dix-Carneiro and Kovak, 2015; Ramondo, Rodríguez-Clare, and Saborío-Rodríguez, 2016; Bakker, 2018). Other work highlights the potential importance of within-country trade costs on international trade (Rauch, 1991; Coşar and Fajgelbaum, 2016), and a large theoretical literature on international specialization arising from agglomeration, initiated by Krugman (1991), has given rise to studies of the stylized interaction between agglomeration and more traditional sources of comparative advantage (Van Marrewijk et al., 1997; Ricci, 1999; Pflüger and Tabuchi, 2016).

Finally, our empirical framework builds on existing analyses of sources of comparative advantage across countries; recent studies that rely on a similar framework include Nunn (2007), Costinot (2009), Chor (2010), Bombardini, Gallipoli, and Pupato (2012), and Cingano and Pinotti (2016).

The chapter is organized as follows. Section 2.2 provides a simple formalization of our hypothesis that comparative advantage across countries stems, in part, from the distribution of population within countries. Section 2.3 describes the data used in the empirical analysis. Section 2.4 presents our main results and Section 2.5 concludes.

2.2 Theoretical Framework

We present a model that illustrates how within-country heterogeneity in productivity can affect a country's pattern of exports across industries. We emphasize how two key ingredients—sector-neutral productivity heterogeneity across a country's locations and differential returns to agglomeration across industries—can produce patterns of specialization both within and across countries. The theoretical results guide our estimation of the key components of our empirical analysis.

2.2.1 Environment: the closed economy

We study an economy in which countries exhibit domestic heterogeneity across inhabited locations, or "cities." A country i is a continuum of cities, indexed by $c \in C_i$, with innate productivity A_c, land area B_c, and equilibrium population L_c. The country's total population is $\bar{L}$; workers are mobile across regions within a country, but not across borders. The economy consists of J tradable sectors indexed by $j = 1, ..., S$, as well as a non-tradable good specific to each city, "housing" (H_c). Tradable goods can be shipped from city c to city d, by paying iceberg trade costs $\tau_{c,d} \geq 1$.

Consumption

Workers in city c earn wage w_c. They derive utility U_c from the consumption of housing and a basket of tradable sectors:

$$U_c(h_c, c_{j=1,...,J}) = (\frac{h_c}{\beta})^{\beta} (\frac{\Pi_{j=1}^{S} (\frac{c_j}{\alpha_j})^{\alpha_j}}{1-\beta})^{1-\beta}$$

where h_c is the worker's housing and c_j, total consumption of sector j, is a CES aggregate of a continuum of varieties indexed by ω, with elasticity σ.

With free within-country trade, the price level in each sector j is common across cities and equal to: $p_j = \left(\int_0^1 p_j(\omega)^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}}$, and the aggregate tradable price level in the country P is $P = \Pi_{j=1}^{J} p_j^{\alpha_j}$. We assume that $\sigma > 1$, so that within each sector, varieties are substitutes. In a spatial equilibrium, utility for a worker with income Y_c is equalized acros cities:

$$V_c = \frac{Y_c}{P^{1-\beta} p_{hc}^{\beta}} = \bar{U} \forall c \tag{2.1}$$

The supply of land in location c is fixed at B_c, the key local congestion force in the model.[10] Equalizing housing supply and demand yields equilibrium housing prices in each city:

$$p_{Hc}^{\frac{1}{\xi}} = \beta \frac{L_c Y_c}{B_c P^{\frac{\xi-1}{\xi}}} \tag{2.2}$$

All Ricardian rents accruing to local landowners are fully taxed by the city government and rebated to resident workers as lump-sum transfers T_c, as in Helpman (1998). Thus, the disposable income Y_c of a worker in city c is proportional to wage income w_c: $Y_c = w_c + T_c = \frac{w_c}{1-\beta\xi}$. Under the spatial equilibrium condition (2.1), city-specific wages are:

$$w_c = P(1-\beta\xi) \bar{U}^{\frac{1}{1-\beta\xi}} \beta^{\frac{\beta\xi}{1-\beta\xi}} \frac{L_c}{B_c}^{\frac{\beta\xi}{1-\beta\xi}} \propto P \times D_c^{\frac{\beta\xi}{1-\beta\xi}}$$

where D_c is the population density of city c.

Production

To study the impact of density on industrial geography and trade, we turn to the supply side of the economy.For simplicity, labor $L_{jc}(\omega)$ is the only input to production. In each industry j, the output of variety ω in city c, $Q_{jc}(\omega)$, is given by:

$$Q_{jc}(\omega) = \tilde{A}_{jc} L_{jc}(\omega)$$

[10] As in Gaubert (2018), atomistic landowners in city c own an amount γ of local land, and produce housing using land and tradable goods, according to the production function: $H_c(\gamma) = \gamma^{\xi} (\frac{X_{hc}(\gamma)}{1-\xi})^{1-\xi}$. For simplicity, we assume that they divide spending on final goods used as inputs in housing production across the S sectors in the same manner as workers; alternatively, one could model the other input into housing production as migrant labor living at zero cost on rural land and only consuming the final good. The details are given in Appendix 2.B.

Each city draws a Ricardian productivity parameter in each variety of good j in location c, $\tilde{A}_{jc}$, from a Fréchet distribution, with cumulative distribution function:[11]

$$\Pr(\tilde{A}_{cj}(\omega) \leq \tilde{A}) = F_{jc}(\tilde{A}) = \exp(-(\frac{\tilde{A}}{A_{jc}})^{\theta})$$

The unit cost of production for variety ω in sector j and location c is then $\frac{w_c}{\tilde{A}_{cj}}$.

Here we introduce the key assumption of the model, which allows us to isolate our channel of interest: the relationship between population distribution and comparative advantage. We assume that the scale of a sector's productivity in city c depends on (i) the city's exogenous sector neutral productivity term A_c, (ii) the city's equilibrium population density D_c, and (iii) the extent to which each sector benefits from local density, $\tilde{\eta}_j$. In particular, we let: $A_{jc} = A_c D_c^{\tilde{\eta}_j}$.

The sector-specific "density elasticity," $\tilde{\eta}_j$, mediates the relationship between density and sector-specific productivity. Variation in $\tilde{\eta}_j$ across sectors—the extent to which each sector benefits from local agglomeration—will be central to our empirical analysis, and is the key modeling assumption. The idea that industries benefit differentially from urban density has been argued in prior work (e.g. Nakamura, 1985; Rosenthal and Strange, 2004; Faggio, Silva, and Strange, 2017) and corroborated by our estimates in Section 2.3.[12]

Trade across cities

If we make the (admittedly strong) assumption that trade costs are zero within country, cost minimization by consumers in any location d then implies that the share of spending on varieties from location c in sector j must be equal for any locations d in the same country:

$$\pi_{dcj} = \pi_{cj} = \frac{p_{cj} X_{dcj}}{X_{dj}} = \frac{(A_c D_c^{\tilde{\eta}_j})^{\theta} w_c^{-\theta}}{\sum_{c'} (A_{c'} D_{c'}^{\tilde{\eta}_j})^{\theta} w_{c'}^{-\theta}} \tag{2.3}$$

[11] We assume the distribution has shape parameter $\theta > \sigma - 1.\theta$, which governs the variance across varieties, is assumed constant across both locations and sectors. As is traditional in supply-driven models of specialization, $\theta > \sigma - 1$ ensures that the CES price index for each sector is well defined.

[12] We remain agnostic here about the specific source of sector-specific density affinity; in section 2.4.5, we explore potential determinants of $\tilde{\eta}_j$.

where π_{dcj} denotes spending in city d on goods in sector j produced in city c.[13]

Equilibrium

Goods market clearing In the equilibrium of the closed domestic economy, the wage bill in each sector j and city c equals total spending on goods produced in sector j in city c.[14] This generates the tradable goods market clearing condition:

$$w_c L_{jc} = \alpha_j \frac{(A_c D_c^{\tilde{\eta}_j})^\theta w_c^{-\theta}}{\sum_{c'} (A_{c'} D_{c'}^{\tilde{\eta}_j})^\theta w_{c'}^{-\theta}} \sum_d w_d L_d \tag{2.4}$$

In the absence of within-country trade costs, the price index for good j is independent of the location where it is consumed and is proportional to:[15]

$$p_j \propto \big[\sum_{c'} (A_{c'} D_{c'}^{\tilde{\eta}_j})^\theta w_{c'}^{-\theta}\big]^{-\frac{1}{\theta}} \propto \big[\sum_{c'} (A_{c'} D_{c'}^{\tilde{\eta}_j - \frac{\beta\xi}{1-\beta\xi}})^\theta\big]^{-\frac{1}{\theta}} \tag{2.5}$$

Trade balance requires that tradable spending from all locations on all goods produced in location c is equivalent to the total wage bill in location c:

$$w_c L_c = \sum_j \sum_d \pi_{dcj} \alpha_j (1-\beta\xi) Y_d L_d = \sum_j \alpha_j \pi_{cj} \sum_d w_d L_d = \sum_d w_d L_d \sum_j \alpha_j \pi_{cj} \tag{2.6}$$

Moreover, the housing market must clear in every location, as in Equation (2.2).

Labor market clearing The ratio of labor allocated to sectors j and j' in each city c is given by:

$$\frac{L_{jc}}{L_{j'c}} = \frac{\alpha_j}{\alpha_{j'}} (\frac{p_j}{p_{j'}})^\theta D_c^{\theta(\tilde{\eta}_j - \tilde{\eta}_{j'})} \tag{2.7}$$

[13] This expression is derived in Appendix 2.B and relies on standard Eaton-Kortum algebra similar to Costinot, Donaldson, and Komunjer (2011) and Michaels, Rauch, and Redding (2013). Given the unbounded nature of the Fréchet distribution, the production structure does not lead to the full specialization of cities in the production of some sectors, which would make the exposition more involved by inducing censoring at the bottom of the sector-city employment density without adding substantial insight in the model.

[14] Note that sector j spending coming from location d is equal to the sum of consumer spending ($\alpha_j(1-\beta)Y_d L_d \pi_{jc}$) and intermediate spending by housing producers ($\alpha_j \beta(1-\xi) Y_d L_d \pi_{jc}$), so that total spending in d on j goods produced in c is $\alpha_j(1-\beta\xi)Y_d L_d \pi_{jc} = \alpha_j w_d L_d \pi_{jc}$.

[15] The proportionality coefficients are independent of the sector and city, since θ is assumed constant

Total population in a city equals the sum of employment across tradable sectors:

$$\sum_j L_{jc} = L_c \tag{2.8}$$

The labor market clears for the country as a whole:

$$\sum_c L_c = \sum_c \sum_j L_{jc} = \bar{L} \tag{2.9}$$

We can now define the equilibrium of the domestic economy.

Definition 2.2.1 (Equilibrium)**.** *An equilibrium in the closed economy is defined as an allocation of labor L_{jc} across cities and sectors such that utility is equalized across sites; trade shares satisfy (2.3); labor allocations satisfy (2.7), (2.8) and (2.9); wages satisfy (2.6) and (2.4); tradable prices satisfy (2.5); and housing prices satisfy (2.2).*

2.2.2 Implications

Within-Country Specialization

We now investigate the domestic sorting of production generated by the model. Double-differencing spending shares (2.3) from any location d across two goods j and j' and locations c and c':

$$\left(\frac{\pi_{jc}}{\pi_{j'c}}\right) / \left(\frac{\pi_{jc'}}{\pi_{j'c'}}\right) = \frac{D_c}{D_{c'}}^{\theta(\tilde{\eta}_j - \tilde{\eta}_{j'})} \tag{2.10}$$

While the absolute unit cost of production is increasing in density D_c due to the need to compensate workers with higher nominal wages, as D_c increases costs increase relatively less fast in sectors with higher $\tilde{\eta}_j$. Denser cities thus have a comparative advantage in sectors that benefit more from agglomeration.[16] Immediately, this implies:

[16]Introducing decreasing returns at the establishment level, for example related to the use of a fixed factor in production such as management skill or land, would make these cross-cities, within-country comparative advantage results hold in terms of the number of establishments as well, consistent with our empirical results in section 2.4.

Lemma 2.2.1. *The share of the labor force employed in higher $\tilde{\eta}_j$ sectors is relatively larger in denser cities:*

$$(\frac{L_{jc}}{L_{j'c}})/(\frac{L_{jc'}}{L_{j'c'}}) = (\frac{w_c L_{jc}}{w_c L_{j'c}})/(\frac{w_{c'} L_{jc'}}{w_{c'} L_{j'c'}}) = (\frac{\pi_{jc}}{\pi_{j'c}})/(\frac{\pi_{jc'}}{\pi_{j'c'}}) = \frac{D_c}{D_{c'}}^{\theta(\tilde{\eta}_j - \tilde{\eta}_{j'})} \tag{2.11}$$

We use Equation (2.11) in our empirical analysis to estimate the $\tilde{\eta}_j$ for each sector (see Section 2.3.3).

Cross-Country Specialization

The following proposition clarifies the implications of the model for country-by-sector level prices in autarky:

Proposition 2.2.1. *The relative price level of two sectors j and j' in the Home country in autarky is:*

$$\log(\frac{p_j}{p_{j'}}) = (\tilde{\eta}_{j'} - \tilde{\eta}_j) \sum_c \omega_{jj',c} \ln(D_c) \tag{2.12}$$

where $\omega_{jj',c}$ are bilateral Sato-Vartia weights (Sato, 1976; Vartia, 1976) across any two goods j and j' in city c, computed from the export shares:

$$\omega_{jj',c} = \Big(\frac{\pi_{cj} - \pi_{cj'}}{\log(\pi_{cj}) - \log(\pi_{cj'})}\Big) \Big/ \Big(\sum_d \frac{\pi_{dj} - \pi_{dj'}}{\log(\pi_{dj}) - \log(\pi_{dj'})}\Big)$$

Proof. See Appendix 2.B. □

Conditional on a fixed distribution of location-level population, the closed economy price index in sector j relative to j' is lower when $\tilde{\eta}_j > \tilde{\eta}_{j'}$. Stronger agglomeration forces in a sector increase productivity in all cities, and lower equilibrium prices for any distribution of density. Moreover, we have the following corollary:

Corollary 2.2.1. *Conditional on the vector of A_c's and wages, a more dispersed distribution of D_c across places – defined as second-order stochastic dominance of the D_c distribution – lowers the price index by more for high $\tilde{\eta}_j$ sectors than for lower $\tilde{\eta}_{j'}$ sectors.*[17]

[17] This follows immediately from Proposition 2.2.1, since the log is concave and $\theta > 0$. As in Proposition 1 in Redding and Weinstein (2020), this results from substitutability across suppliers (note we assumed $\theta > \sigma - 1 > 0$), making the price index log sub-modular in $\tilde{\eta}_j$ and $D_{c'}$.

A more dispersed population implies relatively more variation in sourcing prices across producing locations for higher $\tilde{\eta}_j$ sectors. Substitution across sourcing cities implies lower relative price indices for more "density-loving" sectors in countries with a more dispersed population. This sub-modularity property of price indices in $\tilde{\eta}_j$ and D_c is at the core of comparative advantage of countries in our global economy.

Comparative Advantage To illustrate the implications of the model for patterns of exports under international trade, we aggregate trade flows at the country level. As in Ramondo, Rodríguez-Clare, and Saborío-Rodríguez (2016), we study the special case of N countries, indexed by i, each composed of a set of regions $c \in C_i$, trading S goods indexed by j. We continue to assume that iceberg trade costs are zero across two regions within any country; we also assume trade costs are symmetric and constant across any two regions in two different countries.

To make the results as stark as possible, we assume all countries have the same total population $\bar{L} = L_i$ and the same land area $\int_{c\in C_i} B_c = \int_{c\in C_{i'}} B_c$. We let $B_c = 1$ in each city, so that we simplify the model to the case where $L_c = D_c$. We define X_{inj} as exports from country i to country n in industry j, $\tilde{w}_{ij} = \frac{\sum_{c\in C_i} w_c L_{jc}}{\sum_{c\in C_i} L_{jc}}$ as the average wage in sector j in country i, and M_i as country i's aggregate wage bill, $M_i = \tilde{w}_i L_i = \sum_j w_{ij} L_{ij}$. We can then state the following aggregation result:

Proposition 2.2.2. *Exports of sector j from country i to country n satisfy the following aggregation results*

$$X_{inj} = \alpha_j M_n \frac{T_{ij} \tilde{w}_{ij}^{-\theta} \tau_{ni}^{-\theta}}{\sum_s T_{sj} \tilde{w}_{sj}^{-\theta} \tau_{ns}^{-\theta}}$$

where the country level productivity parameter is:

$$T_{ij} = \Big(\sum_{c\in C_i} (A_c D_c^{\tilde{\eta}_j})^{\frac{\theta}{1+\theta}} (\frac{L_{jc}}{L_{ji}})^{\frac{\theta}{1+\theta}} \Big)^{1+\theta}$$

Moreover, the aggregate wage bill can be expressed as:

$$M_i = \sum_j \tilde{w}_{ij} L_{ij} = \sum_j \Delta_{ij} L_{ij}^{\frac{\theta}{1+\theta}} T_{ij}^{\frac{1}{1+\theta}}$$

where Δ_{ij}, country i's market access in sector j, solves the system of $N \times S$ equations:

$$\Delta_{ij} = \left[\alpha_j \frac{\sum_n M_n \tau_{in}^{-\theta}}{\sum_s \tau_{is}^{-\theta} \Delta_{sj}^{-\theta} L_{sj}^{\frac{1}{1+\theta}} T_{sj}^{\frac{1}{1+\theta}}}\right]^{\frac{1}{1+\theta}}$$

Proof. See Appendix 2.B. □

Country-by-sector productivity T_{ij} is relatively higher for high $\tilde{\eta}_j$ goods in countries with a more concentrated population (and thus, all else equal, for countries with more variance in sector-neutral productive amenities A_c). Even though we all countries have the same total population, the within-country population distribution drives patterns of cross-country trade.

Two-Country Case To build the intuition behind this result, we focus on the the case of two countries, Home and Foreign. First, suppose that Home and Foreign have identical distributions of amenities, A_c and A_c^*. Then there will be cross-city trade both within and across countries, but there will be no apparent pattern of inter-industry trade at the country level. Next, assume the distribution of *sector-neutral* productivity across cities is *more even* in the Foreign country than at Home. By "more even", we mean that the distribution of Foreign productivity is a "utility-preserving spread," an extension of the "mean-preserving spread" concept defined as:

Definition 2.2.2. *G is a "utility-preserving spread" of G^* if in the closed economy, welfare is the same at Home and in Foreign, $\bar{U} = \bar{U}^*$, but the variance of A_c is higher than the variance of A_c^*.*[18]

This implies, from Equation (2.13), that the distribution of population at Home second-order stochastically dominates the distribution in Foreign; the Generalized Lorenz Curve of population in the Foreign economy lies strictly above the Lorenz curve at Home. By Proposition 2.2.1, the relative prices of higher $\tilde{\eta}_j$ goods are lower in the closed Home economy than in the closed Foreign economy. Equation (2.11) implies that the relative share of employment of high $\tilde{\eta}_j$ sectors is increasing in density, so in the Home country, relatively more

[18]One can imagine an experiment with two cities, c_1 and c_2, where initially $A_{c_2} > A_{c_1}$. Then a utility-preserving spread could involve lowering A_{c_1} by ϵ, and increasing A_{c_2} by $\alpha\epsilon$, where α is chosen so that $\bar{V}_0 = \bar{V}'(\alpha)$.

workers are active in high $\tilde{\eta}_j$ sectors than in the Foreign country. Aggregating cross-location trade flows to the country level, the Home country will appear to specialize in goods that have a high $\tilde{\eta}_j$'s and import goods with lower $\tilde{\eta}_j$'s. Let the Generalized Lorenz Curve (GLC) of population density be the cumulative distribution function of experienced density.Then, in a two-good setting:

Corollary 2.2.2. *Suppose there are two countries, H and F, and two goods j and j' where $\tilde{\eta}_j > \tilde{\eta}_{j'}$ and $\alpha_j = \alpha_{j'}$. The CDFs of location-specific amenities in H and F are G and G^*. If G is a utility-preserving spread of G^*, then the GLC of population-weighted density in H lies strictly below the Generalized Lorenz Curve of population-weighted density in F. Moreover, H is a net exporter of j and F is a net exporter of j'.*

From Theory to Measurement: Population-Weighted Density

From the equilibrium definition in 2.2.1, the population distribution can be expressed as the labor market clearing (2.9), along with a system of C equations that depend on city-level population-weighted density, city-level population weighted amenities, and a constant term:

$$L_c D_c^{\frac{\beta\xi}{1-\beta\xi}} = \sum_j \alpha_j \frac{(A_c D_c^{\tilde{\eta}_j - \frac{\beta\xi}{1-\beta\xi}})^\theta}{\sum_{c'} (A_{c'} D_{c'}^{\tilde{\eta}_j - \frac{\beta\xi}{1-\beta\xi}})^\theta} \sum_d L_d D_d^{\frac{\beta\xi}{1-\beta\xi}} \tag{2.13}$$

There is a unique equilibrium when the maximum sector-level density elasticity ($\tilde{\eta}_{max} = \max_j \tilde{\eta}_j > 0$) is "not too large" relative to the share of land in housing production (ξ); this makes congestion forces strong enough to offset multiple equilibria.[19]

At the country level a greater dispersion of A_c leads to greater equilibrium D_c dispersion. In particular, the population density distribution in an economy with more dispersed A_c is second-order stochastically dominated by the population density distribution in an economy with less dispersed A_c^* (see Appendix 2.B), and we will observe the footprint of productivity dispersion across cities in the dispersion (or concentration) of population. In the special case where total population is held constant, which we ensure in our empirical

[19]The proof is analogous to Redding (2016). For a sufficiently small $\tilde{\eta}_{max}$, a location's density D_c is increasing in its productive amenity A_c, since a higher A_c increases the marginal product of labor in any sector, leading to rising nominal wages, population inflows, and land prices, until utility is again equalized. Agglomeration forces, modeled as positive $\tilde{\eta}_j$'s, reinforce this phenomenon, but do not offset it if they are small enough.

analysis, and B_c and A_c are uncorrelated, greater dispersion in the exogenous A_c's can be mapped directly to greater country-level "population-weighted density":

$$D_i = \int_0^{\max D_c} \frac{L_c^2}{B_c} dH(D_c)$$

which captures the local population density experienced by the average worker in the economy. While, as discussed below, there are several intuitively appealing features of using this as our county-level parameterization of population concentration, the model also indicates that it is the observable consequence of dispersion (or lack thereof) of the primitive productivity distribution. This is the measure we estimate next in Section 2.3, and use as our main measure of population concentration ("density") in Section 2.4.1.

2.3 Measurement

2.3.1 Data Sources

Economic Geography Data on economic activity in the US are collected from the 2016 version of the County Business Patterns (CBP) data set. The CBP contains information on employment, establishment counts, and total payroll in each industry and Core-Based Statistical Area (CBSA). We focus on measures at the NAICS 4-digit level, which are less likely to suffer from suppression.[20] We compile data on a range of industry-level characteristics from the latest available year in the NBER-CES Manufacturing Industry Database, including capital intensity, the labor share, and average wages. We also include data from the American Community Survey to control for the age and gender breakdown of the workforce as well as detailed measures of the educational attainment of the workforce in each industry.

To construct instruments for local density, we also compile data on distance to subterranean bedrock for all US CBSAs. Raster data displaying the distance to bedrock of each 250m grid cell in the US, which we use to construct the instruments, are from the International Soil Reference and Information Centre (ISRIC) *SoilGrid* project.[21]

[20]We verify that our results are not sensitive to imputation when using interpolation techniques to impute missing employment data in the CBP.

[21]See here: https://www.isric.org/explore/soilgrids.

Density Spatial data on global population density are obtained from the *LandScan* Database.[22] These data are calculated by combining existing demographic and census data with remote sensing imagery, and are released as a raster data set composed of one square-kilometer grid cells.[23] The resulting population count is an ambient or average day/night population count. We use the the *LandScan* data to compute state and country-level estimates of population-weighted density. For our instrumental variables analysis, we also rely on new measures of historical population and city size distributions constructed from data sets recently introduced by Reba, Reitsma, and Seto (2016) and Fang and Jawitz (2018).

Trade US State-level international exports from 2016 are collected from the US Census Bureau's USATradeonline database.These data are provided at the NAICS 4-digits level, which is our primary level of analysis across industries. We focus on gross exports flows, as they are the natural counterpart of spending in our theoretical framework. Cross-country trade flows data are obtained from the UN Comtrade Database for all available exporters in 2016, at the HS4 digit level.We map HS4 industries to NAICS-4 industries using the crosswalk developed by Pierce and Schott (2012).

Additional Data To include additional controls in our cross-state and cross-country estimates, we compiled US state-level data on educational attainment, age composition, and worker income from the 2016 American Community Survey estimates. At the country level, we also compiled information on educational attainment, urbanization, GDP per capita, and a range of other country-level characteristics from the World Bank's World Development Indicators and International Monetary Fund's World Economic Outlook databases, and mea-

[22]*LandScan* data can be found here: `https://landscan.ornl.gov` We use the LandScan data product from 2016.

[23]For more information, see here: https://landscan.ornl.gov/documentation. According to *LandScan*: *ORNL's LandScan is the community standard for global population distribution. At approximately 1 km resolution (30 $\times$ 30 degree), LandScan is the finest resolution global population distribution data available and represents an ambient population (average over 24 hours). [...] The LandScan global population distribution models are a multi-layered, dasymetric, spatial modeling approach that is also referred to as a "smart interpolation" technique. In dasymetric mapping, a source layer is converted to a surface and an ancillary data layer is added to the surface with a weighting scheme applied to cells coinciding with identified or derived density level values in the ancillary data. [...] The modeling process uses sub-national level census counts for each country and primary geospatial input or ancillary datasets, including land cover, roads, slope, urban areas, village locations, and high resolution imagery analysis; all of which are key indicators of population distribution. [...] Within each country, the population distribution model calculates a "likelihood" coefficient for each cell and applies the coefficients to the census counts, which are employed as control totals for appropriate areas. The total population for that area is then allocated to each cell proportionally to the calculated population coefficient*

sures of country-level capital stocks from the Penn World Tables.

2.3.2 Estimating State and Country Level Density

For both US states and countries, we compute *population-weighted density* (D_i) as:

$$D_i = \sum_{g \in G(i)} \left(L_g \times \frac{L_g}{\sum_{g' \in G(i)} L_{g'}} \right)$$

where g indexes grid cells and $G(i)$ is the set of grid cells in country (or state) i. L_g is the population, according to *LandScan*, in grid cell i. Since all grid cells are the same size, L_g is also the density of grid cell i. This measure is equivalent to weighting the population density of each grid cell in a country or state by its population, and yields a measure of population density that approximates to the expected experienced density of a person in the state or country.[24]

This is our key state and country-level independent variable of interest. Intuitively, this measure captures the concentration of population within a state or country. For a given total population if people are very concentrated in a few cities this measure will be large whereas if people are is dispersed across many less-dense cities or suburban and rural areas, D_i will be small. Figure 2.3.1 plots the distribution of D_i across US states. While, intuitively, populous and urban states like New York and California have high measures of D_i, so do Massachusetts and Washington; large states like Texas and Florida, with their large but more sprawling cities, are in the middle of the distribution. Figure 2.1.1 (above) had displayed deciles of D_i for each country around the world. There is substantial variation in D_i across countries, both within continents and within income groups.

2.3.3 Estimating Sector-Specific Density Affinity

Using industry-by-city level data from the US County Business Patterns (CBP), we estimate the agglomeration elasticity of each tradable manufacturing sector. Because our focus is cross-country trade, and manufactured goods account for the bulk of international exports,

[24] See Wilson (2012) for a justification of the use of population-weighted density by the United States Census Bureau.

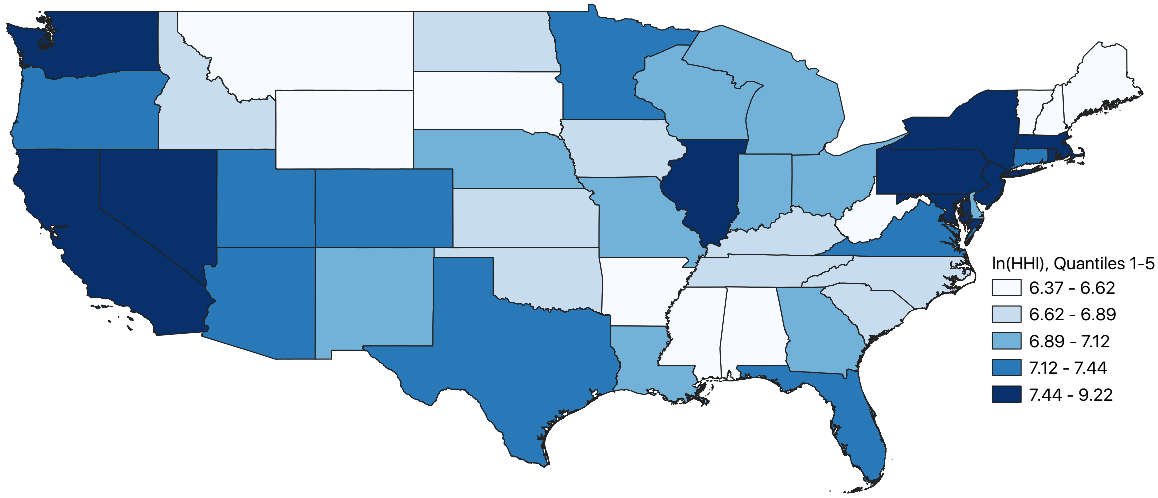

Figure 2.3.1: **Population weighted population density across US states.**

The figure is a map in which US states are color-coded based on their population-weighted density quintile. Darker shaded states have higher population-weighted density.

we emphasize the existence of substantial within-manufacturing differences in density affinity.

We compute a "density-elasticity" for each industry by estimating the following empirical analog of the the model's Equation (2.11):

$$y_{cj} = \alpha_c + \gamma_j + \sum_j \eta_j \cdot \left(\ln D_c \cdot \mathbb{I}_j\right) + \epsilon_{cj} \tag{2.14}$$

where c indexes cities and j indexes sectors. y_{cj} is the (log of the) number of employees, number of establishments, or first quarter aggregate payroll in industry j and location (city) c. α_c and γ_j are city and sector fixed-effects, respectively. D_c is population density at the level of the Core Based Statistical Area (CBSA) and $\mathbb{I}_j$ is an indicator that equals one for sector j. The coefficients of interest are the density elasticities, η_j, the key source of industry-level variation in the model. These elasticities capture the extent to which each industry tends to be more or less represented in denser locations.

We first estimate Equation (2.14) using OLS and report the ten sectors with the highest and lowest density elasticities in Table 2.3.1a. Since CBSA-level density is likely correlated with a range of other city-level characteristics that might affect industry sorting, it is difficult to interpret the purely correlational estimates. To circumvent this issue, we construct an instrument for CBSA-level density in order to estimate the causal effect of a marginal change in CBSA-level density on industry-specific production. Subterranean geology affects ease of vertical construction, and hence potential population density, but is unlikely to independently affect other city-level characteristics. Our instrument is the (log of the) average distance of each CBSA to subterranean bedrock. Lower distance to bedrock in a location eases the land constraint, and can be interpreted as increasing the available share of land B_c in our theoretical framework; construction often requires a foundation in bedrock and is more difficult when bedrock is deep (e.g. Schuberth, 1968; Landau and Condit, 1999).[25] By exogenously shifting density, we estimate the response of industry specialization to density alone, capturing the causal effect of a marginal change in city-level density on industry-level production.

The correlation between CBSA-level density and the log of the distance to bedrock is shown in Figure 2.3.2. The correlation coefficient is highly statistically significant (t-statistic = 8.07) suggesting that, consistent with the mechanical impact of distance to bedrock on construction, CBSA-level variation in subterranian bedrock systematically shifts equilibrium population density. The necessary identification assumption is that distance to subterrenian bedrock only affects industry sorting through its impact on ease of construction and hence population density.[26]

We then estimate $\eta)j$ for each sector using IV-2SLS, and the interaction between industry-level indicators ($\mathbb{I}_j$) and (log of) distance to bedrock as the instruments. Industries with the highest and lowest IV estimates of η_j are listed in Table 2.3.1b. While many of these sectors are intuitive and commonly associated with production in dense cities, in the case of the top sectors, or production away from large cities, in the case of the bottom sectors, they also do not map clearly onto common determinants of comparative advantage. The top of our list features both industries that are skill-intensive (e.g. Semi-conductor and Other Electronic Component Manufacturing) and industries that are not skill-intensive (e.g. Beverege

[25]Recent research has suggested the use of underlying geologic characteristics to provide exogenous sources of variation in land supply availability and estimate its economic effects (Rosenthal and Strange, 2008; Saiz, 2010; Duranton and Turner, 2018) However, existing research has focused on within-city variation in geological features to instrument for urban shape, rather than variation across metropolitan areas.

[26]While this assumption seems likely, we also verify that the results are similar after controlling for other ground and soil characteristics (e.g. characteristics of soil content, agricultural suitability, etc.). These estimates and their possible parameterizations are available upon request.

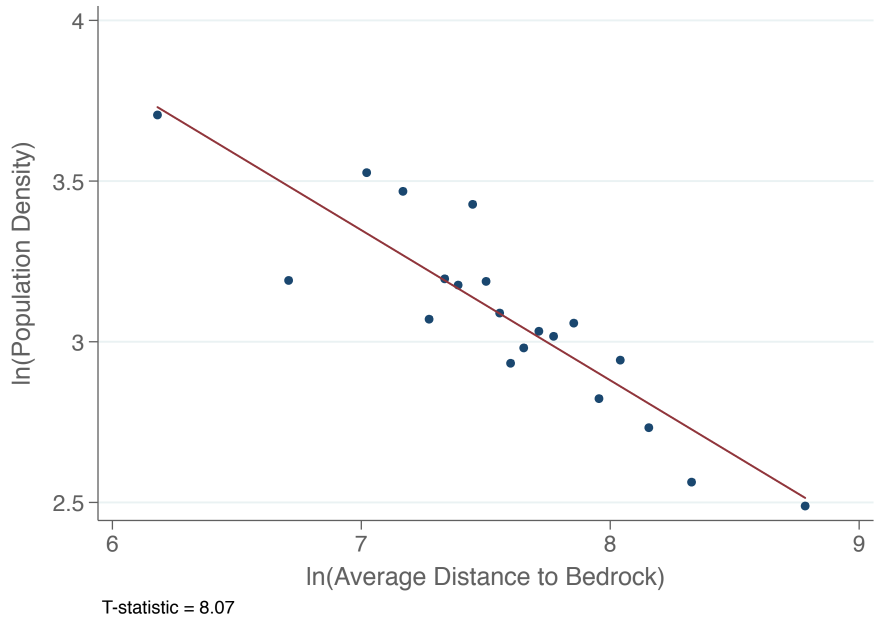

Figure 2.3.2: **Distance to Bedrock and Population Density.**

The figure is a binned scatter plot. It reports the correlation between log of distance to bedrock and log of population density at the CBSA level. The t-statistic is 8.07.

Manufacturing). The same is true for capital intensity.[27]

Figure 2.3.3 shows the distribution of establishments in the top and bottom ten sectors listed in Table 2.3.1b across the US. For each CBSA c and sector j, we compute:

$$\text{Representation}_{cj} = \left(\frac{\sum_{j\in \text{T,B}} \text{Establishments}_{cj}}{\sum_j \text{Establishments}_{cj}}\right) \Big/ \left(\frac{\sum_c \sum_{j\in \text{T,B}} \text{Establishments}_{cj}}{\sum_c \sum_j \text{Establishments}_{cj}}\right)$$

where T and B are the set of ten highest and lowest η_j sectors respectively. This normalization captures the over- or under-representation of top or bottom sectors in city c by normalizing the share of city c manufacturing establishments that belong to $j \in T/B$ by the overall share of manufacturing establishments that belong to $j \in T/B$ in the US.

Figure 2.3.3a shows the geographic distribution of low-η_j sectors; they are disproportionately located in Upper Midwest and Central and Northern Plains regions (purple-shaded regions). High-η_j sectors, displayed in Figure 2.3.3b, are disproportionately located on the

[27]Moreover, motor vehicle manufacturing, for example, the top of Nunn (2007)'s list of contract intensive industries, but are at opposite ends of our list. The same is true of Manufacturing and Reproducing Magnetic and Optical Media.

Table 2.3.1: The Ten Most and Least Density Elastic Industries: OLS and IV Estimates

(1)	(2)	(3)	(4)	(5)	(6)
TOP TEN			BOTTOM TEN		
Elasticity: OLS Estimate	NAICS Code	Industry Name	Elasticity: OLS Estimate	NAICS Code	Industry Name
1.841537	3222	Converted Paper Product Manufacturing	0.5473573	3117	Seafood Product Preparation and Packaging
1.708105	3345	Navigational, Measuring, Electromedical, and Control Instruments	0.5272605	3131	Fiber, Yarn, and Thread Mills
1.702192	3261	Plastics Product Manufacturing	0.5059118	3112	Grain and Oilseed Milling
1.641616	3344	Semiconductor and Other Electronic Component Manufacturing	0.4773871	3365	Railroad Rolling Stock Manufacturing
1.632981	3363	Motor Vehicle Parts Manufacturing	0.4667493	3162	Footwear Manufacturing
1.531617	3339	Other General Purpose Machinery Manufacturing	0.4558011	3361	Motor Vehicle Manufacturing
1.520556	3342	Communications Equipment Manufacturing	0.4041506	3221	Pulp, Paper, and Paperboard Mills
1.508072	3321	Forging and Stamping	0.4014953	3161	Leather and Hide Tanning and Finishing
1.493721	3255	Paint, Coating, and Adhesive Manufacturing	0.3166769	3211	Sawmills and Wood Preservation
1.487333	3353	Electrical Equipment Manufacturing	-0.0162934	3122	Tobacco Manufacturing

Notes: The density elasticity measure is estimated by OLS.

(a) OLS Estimates

(1)	(2)	(3)	(4)	(5)	(6)
TOP TEN			BOTTOM TEN		
Elasticity: IV Estimate	NAICS Code	Industry Name	Elasticity: IV Estimate	NAICS Code	Industry Name
1.524013	3117	Seafood Product Preparation and Packaging	-0.1235644	3361	Motor Vehicle Manufacturing
1.271837	3151	Apparel Knitting Mills	-0.1605782	3331	Agriculture, Construction, and Mining Machinery Manufacturing
1.226573	3342	Communications Equipment Manufacturing	-0.1798471	3112	Grain and Oilseed Milling
1.197432	3121	Beverage Manufacturing	-0.2356362	3325	Hardware Manufacturing
1.165016	3219	Other Wood Product Manufacturing	-0.2846771	3221	Pulp, Paper, and Paperboard Mills
1.147163	3132	Fabric Mills	-0.3778235	3339	Other General Purpose Machinery Manufacturing
1.011703	3371	Household and Institutional Furniture and Kitchen Cabinet Manufacturing	-0.4019563	3111	Animal Food Manufacturing
1.006027	3344	Semiconductor and Other Electronic Component Manufacturing	-0.4702834	3274	Lime and Gypsum Product Manufacturing
0.9815783	3113	Sugar and Confectionery Product Manufacturing	-0.5856257	3114	Fruit and Vegetable Preserving and Specialty Food Manufacturing
0.8887671	3211	Sawmills and Wood Preservation	-0.6302103	3346	Manufacturing and Reproducing Magnetic and Optical Media

Notes: The density elasticity measure is estimated by IV-2SLS.

(b) IV Estimates

East and West coasts, as well as in cities in Texas and parts of the Midwest. There is significant variation within regions and states as well. Indeed, almost all states have locations in which both high and low η_j sectors are disproportionately produced.

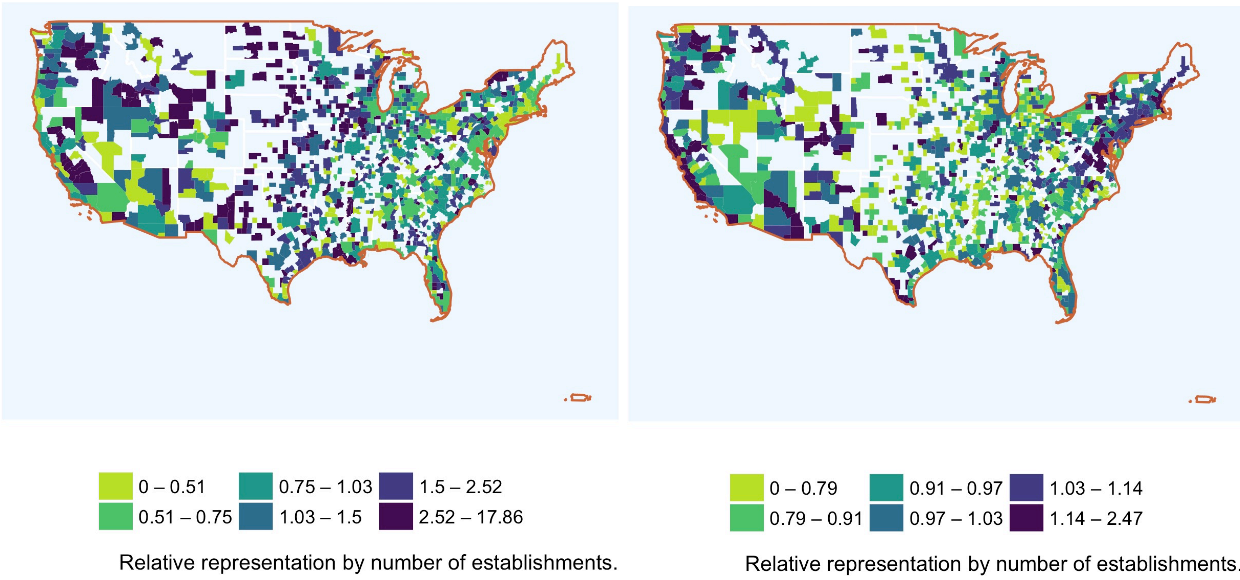

(a) Low η_j Sectors

(b) High η_j Sectors

Figure 2.3.3: **Representation of Low- and High- η_j Sectors Across US Cities.**

Both (a) and (b) are US CBSA-level maps. (a) displays the relative representation of low-η_j sectors, the ten sectors with the lowest first principal component of our six density elasticity estimates. (b) displays the relative representation of high-η_j sectors, the ten sectors with the lowest first principal component of our six density elasticity estimates. These sectors are listed in Table 2.3.1

2.4 Empirical Results: Population Distribution and the Pattern of Trade

2.4.1 Estimation Framework

We now examine the impact of within-country population distribution on patterns of trade. We investigate whether population-weighted density, D_i, is a systematic source of comparative advantage. Our main empirical estimating equation is:

$$y_{ij} = \alpha_i + \gamma_j + \beta \cdot \eta_j^{IV} \cdot \ln(D_i) + X'_{ij}\Gamma + e_{ij} \tag{2.15}$$

where i indexes states or countries and j indexes sectors. The unit of observation is a country (or state)-by-sector pair. The dependent variable is total exports in sector j from state or country i. The independent variable of interest is an interaction term between (i) IV estimates of sector-level density affinity (η_j^{IV}) and (ii) log of state or country-level population weighted density ($\ln(D_i)$). The density affinity of all NAICS-4 sectors were estimated using Equation (2.14) and the instrumental variables strategy outlined in Section 2.3.3. All specifications include sector and state or country fixed effects. We will also include a range of controls that vary at the state-by-sector or country-by-sector level (X'_{ij}); these vary across specifications to probe the sensitivity of our estiamtes. Following Silva and Tenreyro (2006), we use the Poisson pseudo-maximum likelihood (PPML) estimator as our baseline specification, but show throughout that results are similar using OLS and a log-transformed dependent variable.[28]

The coefficient of interest is β. If $\beta > 0$, it implies that countries with greater population-weighted density have a revealed comparative advantage in "density-loving" sectors. This framework follows the regression-based index of comparative advantage summarized in French (2017), as used, among others, by Nunn (2007) or Bombardini, Gallipoli, and Pupato (2012). In Section 2.4.4 (below) we propose an instrumental variables strategy that exploits variation in historical population and city size distributions as shifters of modern population density.

2.4.2 US State-Level Estimates

The over-representation of some manufacturing sectors in dense areas in the United States might stem from either local supply or local demand conditions. Our hypothesis focuses on the supply side, by suggesting that denser cities are relatively more efficient in the production of "density-loving" industries. If this is the case, dense areas within the US should not only attract relatively more employment and production in these industries, but also export significantly more of them internationally. Moreover, while many models of international trade consider the entire US as a single "point," different parts of the US specialize in vastly different industries (see e.g. Irwin (2017) for a long-term perspective). Thus, as a first test of our hypothesis that regions with greater population-weighted density specialize in the export of density-loving industries, we estimate Equation (2.15) at the US state level.[29]

[28] As shown by Fally (2015), the Poisson pseudo-maximum likelihood estimation method has the additional benefit of ensuring that predicted trade flows satisfy the "adding up" constraint implicit in gravity models of trade.

[29] While some recent studies have attempted to estimate export data at the metropolitan level (see e.g. the database constructed by Tomer and Kane (2014)), most trade flows data are still collected at a broader level of

Table 2.4.1: State-Level Trade, Baseline Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Dependent Variable is Total Exports from the State-Sector						
Strategy for estimation of density affinity:	η_j computed using industry-level employment					η_j computed using industry-level number of establishments	
	Panel A: Outcome Variable is Total Exports (Thousands), PML Model						
D_i x η_j	0.612***	0.539***	0.563***	0.437***	0.538***	3.508***	3.241***
	(0.145)	(0.117)	(0.201)	(0.0917)	(0.199)	(0.541)	(0.660)
	Panel B: Outcome Variable is log(Exports), OLS Model						
D_c x η_j	0.146*	0.129*	0.142*	0.120*	0.124	0.864**	0.839**
	(0.0734)	(0.0725)	(0.0738)	(0.0685)	(0.0793)	(0.358)	(0.363)
R-squared	0.756	0.758	0.757	0.758	0.760	0.756	0.760
Factor Intensity Controls	No	Yes	No	No	Yes	No	Yes
State Level Controls	No	No	Yes	No	Yes	No	Yes
Industry Level Controls	No	No	No	Yes	Yes	No	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
States	50	50	50	50	50	50	50
Observations	4,182	4,132	4,182	4,132	4,132	4,182	4,132

Notes: The unit of observation is a state-by-sector pair. The coefficient of interest is the coefficient on an interaction between state-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment in columns 1-5 and estalishments in columns 6-7. Panel A reports Poisson pseudo-maximum likelihood estimates while Panel B reports OLS estimates. All specifications include state and sector fixed effects, along with other controls listed at the bottom of each column. Standard errors, clustered at the state level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively

These estimates are reported in Table 2.4.1. Panel A reports Poisson maximum likelihood estimates while Panel B reports OLS estimates with log of exports as the outcome variable. Across specifications, we find that the coefficient of interest is positive and statistically significant, suggesting that US states with greater population-weighted density have a comparative advantage in density-loving industries. Column 1 presents the coefficient of interest when only $\eta_j^{IV} \times \ln(D_i)$—the interaction between state-level population weighted density and industry-level density affinity—is included on the right hand side (along with state and industry fixed effects). The remaining specifications investigate the robustness of this baseline result to the inclusion of additional controls.

In order to investigate whether the results are driven by state-level differences in education and comparative advantage in high-skill industries, in column 2 we include a series of interactions between state-level educational attainment and sector-level skill demand. In

aggregation. The smallest level of consistent and exhaustive trade reporting in the United States is the state.

particular, we separately interact the share of people in each state who have achieved a (i) high school degree, (ii) a bachelors degree, and (iii) a graduate degree, with the share of people employed in each sector (i) that have a high school degree or (ii) that have at least a college degree. The inclusion of these six interactions has little effect on our coefficient of interest.

In column 3, we control for a series of state-level variables interacted η_j^{IV} in order to investigate whether the baseline result is driven by some omitted state-level characteristic. These controls include (log of) the median household income; (log of) state-level population; the share of inhabitants with high school, bachelor, and graduate degree; and the share of young people, aged 18-30. It is possible, for example, that denser states are also just wealthier and that this drives the baseline estimate. However, the coefficient of interest remains very similar after including these controls.

In order to address the potential for omitted industry-level characteristics, in column 4 we control for a series of industry-level characteristics interacted with $\ln(D_i)$. These covariates, computed for each manufacturing industry in the US, are the value of installed capital per worker, (log of) the average employee compensation, the share of workers with at least a college degree, the average age of employees, and the gender breakdown of employment. In column 5 we include all 17 controls mentioned thus far and again, the coefficient of interest remains very similar. It does, however, lose statistical significance in Panel B when we use an OLS regression model and log of exports as the outcome variable; this is driven by a larger standard error rather than a decline in coefficient magnitude.

In columns 6-7 we repeat the specifications from columns 1 and 5—the specifications without any controls and the specification with all controls—and construct the "density affinity" measure using industry-level establishment data rather than employment data. The number of establishments is a potentially less noisy measure of industry-level production across space than employment, and moreover is never suppressed in the CBP data. Reassuringly, in both columns 6 and 7 and in both Panels A and B, our coefficient of interest is positive and highly significant. Finally, Table 2.A.1 reports estimates from a series of additional specifications; each reported coefficient in Table 2.A.1 is estimated from a separate regression. The results are very similar if we use the versions of η_j estimated using OLS (instead of IV) and using city-level data on payroll, rather than employment or establishments. Finally, all findings are very similar if we exclude state-industry pairs with zero exports (Table 2.A.2).

This first set of results demonstrates that US states that exhibit a more spatially concentrated population export relatively more in sectors whose production is concentrated in denser metropolitan areas. According to our estimates, a one-standard deviation increase in the density interaction in the fully controlled specification increases the dependent variable by 0.139 standard deviations when computed using the elasticity with respect to employment and 0.295 when computed using the elasticity with respect to establishments.

2.4.3 Country-Level Estimates

We now turn to the main results of the paper: the relationship between density and patterns of cross-country trade. Estimates of (2.15) in which the units of observation are country-industry pairs are reported in Table 2.4.2. Panel A presents Poisson pseudo-maximum likelihood estimates while Panel B reports OLS estimates. The coefficient of interest in a specification without controls is presented in column 1; it is positive and highly significant. Countries with a more concentrated population distribution have a revealed comparative advantage in density-loving sectors.

Columns 2-6 investigate the robustness of the result to the inclusion of a series of controls in order address potential concerns due to omitted variable bias. In column 2, we control for traditional determinants of comparative advantage , including capital and skill intensity (Romalis, 2004).[30] Since data on the country-level capital stock is only available for 90 countries, the sample size of the regression is reduced; nevertheless, the coefficient of interest is almost exactly identical.

In column 3 we control for a series of country-level characteristics interacted with the sector-level density elasticity measure, η_j^{IV}. These are included to account for the fact that population-weighted density is potentially related to other country-level characteristics that may affect comparative advantage. In particular, we control for (the log of) country-level

[30] In particular, we interact country-level capital stock (as drawn from the Penn World Tables) with an industry's average level of capital intensity obtained from the NBER-CES Manufacturing database. We also interact measures of educational attainment at the country level with our estimates of the skill intensity of an industry in US data computed from the share of high school and college attainment of workers in the industry in the American Community Survey data.

Table 2.4.2: Country-Level Trade, Baseline Estimates

	(1)	(2)	(3)	(4)	(5)
	Dependent Variable is Total Exports from the Country-Sector				
	Panel A: Outcome Variable is Total Exports (Thousands), PML Model				
D_i x η_j	0.456***	0.464***	0.757***	0.462***	0.765***
	(0.111)	(0.110)	(0.0849)	(0.0710)	(0.0731)
	Panel B: Outcome Variable is log(Exports), OLS Model				
D_i x η_j	0.104**	0.105**	0.288***	0.122***	0.262***
	(0.0487)	(0.0524)	(0.0645)	(0.0454)	(0.0627)
R-Squared	0.814	0.796	0.793	0.816	0.797
Factor Intensity Controls	No	Yes	No	No	Yes
Country Level Controls	No	No	Yes	No	Yes
Industry Level Controls	No	No	No	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
Countries	134	90	107	134	83
Observations	10,464	7,241	8,542	10,332	6,674

Notes: The unit of observation is a country-by-sector pair. The coefficient of interest is the coefficient on an interaction between country-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment. Panel A reports Poisson pseudo-maximum likelihood estimates while Panel B reports OLS estimates. All specifications include country and sector fixed effects, along with other controls listed at the bottom of each column. Standard errors clustered at the country level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

total population, educational attainment, urbanization, the share of population employed in agriculture, the share of population employed in service production, (log of) per capita GDP (PPP adjusted), and a rule of law index, all interacted with η_j^{IV}. Again, the coefficient of interest is very similar.[31] Further, in Table 2.A.3 we reproduce our findings after including continent-by-industry fixed effects; this specification flexibly controls for differences in *industry-specific* productivity and trade in different parts of the world.

Next, we investigate the robustness of the result to the inclusion of sector-level controls. We control for the same industry-level controls as in Table 2.4.1, interacted with country-level measures of population-weighted density, D_i. Reassuringly, the coefficient of interest is again very similar after the inclusion of these controls.

[31] Moreover, the coefficient of interest is also similar if only individual country-level controls or smaller sets of country-level controls are included on the right-hand side, but to conserve space we do not report these specifications.

In column 5, we include all controls mentioned thus far on the right-hand side of the regression. Due to missing covariates, the sample size is reduced to 83 countries, yet the coefficient of interest remains positive and highly significant, suggesting that our findings are not driven by standard determinants of comparative advantage or other measurable country or industry level covariates. Table 2.A.5 documents that the results are not sensitive to the use of our alternative estimates of η_j, and are also very similar after excluding countries in the bottom 10% of the income and population distributions. The lowest income countries likely also have lower quality data and the smallest or poorest countries might have extreme values of either density or trade values. As in the case of our state-level estimates, the findings are also very similar if we include country-industry pairs with zero exports (Table 2.A.2).

Finally, we investigate the robustness of the findings to alternative potential sources of population data. While our baseline results rely on the Landscan database, other organizations, using slightly different methodologies to account for sparse data in some parts of the world, also produce gridded global population estimates. These are: the Global Human Settlement Layer, the Gridded Population of the World, and the WorldPop Project. We measure country-level population density using each of these data sets and re-estimate our baseline results after computing the independent variable of interest from each data source. These results are presented in Table 2.A.4 and our findings are very similar across population data sources.[32]

These estimates indicate that the distribution of population within countries is a potentially important determinant of comparative advantage and patterns of trade. Our point estimate from column 2, when only factor endowment controls are included, implies that a one standard deviation increase in the density interaction increases the outcome variable by 0.113 standard deviations. This is slightly larger in magnitude than the coefficient on the capital interaction, which implies a standardized beta coefficient of 0.109.[33] In the specification with all controls included, the coefficient of interest increases and implies a beta

[32]We thank Richard Delome for pointing this out to us, and rely on his version of the data sets which can be found here: https://github.com/richarddelome/density_metrics/blob/master/README.md?fbclid=IwAR1KQ1KJB5FeLW45R0HXA63gfET9XT8jS7ecmaQ9h-B7LmPYuJW1ODAdK98.

[33]Reassuringly, our estimates of the magnitudes of comparative advantage due to factor endowments is very similar to Nunn (2007), who estimates a beta coefficient on an analogous capital interaction of 0.105.

coefficient on the density interaction of 0.276.

2.4.4 Endogeneity

This section proposes an instrument for population-weighted density and reports instrumental variable estimates of our baseline specification. The goal of introducing an instrument is to make sure that the baseline results are not driven by reverse causality. That is, it is possible that the composition of a state or country's exports has feedback effects and shapes its economic geography; we would then find a positive coefficient on our density interaction, but it would be incorrect to interpret the relationship as evidence that density is a source of comparative advantage. To rule out the possibility that our results capture the effect of trade on economic geography, we use characteristics of a state or country's *historical* population distribution to construct instruments for the population distribution today. While characteristics of a country's historical population distribution predict its modern population distribution, it seems unlikely that modern patterns of trade, which developed largely after World War II, had a direct effect on the population distribution in in 1900 (e.g. Irwin, 2017).

The ideal instrument for our purposes would be a historical measure of population weighted density, analogous to our contemporary measure. We construct such a measure for each US state using estimates of the historical US population distribution presented in Fang and Jawitz (2018). Fang and Jawitz (2018) combine historical census data with population modeling techniques to construct a spatially explicit distribution of the US population for each decade since 1790.[34] Using this gridded data set, we compute the population weighted density of each US state in 1900 (D_i^{1900}).[35] The first stage estimating equation is thus:

$$(\eta_j^{IV} \cdot \ln(D_i)) = \xi \cdot \eta_j^{IV} \cdot \ln(D_i^{1900}) + \alpha_i + \gamma_j + +X_{ij}'\Gamma + e_{ij} \tag{2.16}$$

where we hypothesize $\xi > 0$ if the historical state-level population distribution is a strong

[34]While the most advanced version of their model also uses socioeconomic characteristics of each region to predict population, we use the "Level 4" version of the model that does not take socioeconomic characteristics into account.

[35]We select the year 1900 for comparability with our country-level IV estimates, which have additional data constraints and are reported below.

predictor of the modern population distribution.

Out state-level IV-2SLS estimate of Equation (2.15), where the first stage estimating equation is (2.16), is presented in column 1 of Table 2.4.3. The coefficient estimate is positive, statistically significant, and similar in magnitude to the OLS estimates, suggesting that our state-level findings are not driven by reverse causality. Moreover, the first stage relationship is also strong; the Kleibergen-Paap first stage F-statistic is 25.159.

While it is possible to estimate the historical population weighted density of each US state, to our knowledge this is not possible at the country level. Therefore, in order to adapt the logic of our identification strategy to the country-level analysis, we also introduce a second set of instruments. We determined the location and population of cities around the world in 1900 using historical data collected by Chandler (1987), and recently digitized by Reba, Reitsma, and Seto (2016).[36] Intuitively, high D_i corresponds to having a high city population concentrated in a relatively small number of cities. For each state and country, we therefore compute the total population across all cities (p_i^{1900}), as well as the inverse number of cities (c_i^{1900}). We include both, as well as their interaction ($p_i^{1900} \cdot c_i^{1900}$), interacted with η_j, as excluded instruments. We expect $p_i^{1900} \cdot c_i^{1900} \cdot \eta_j$ to be positively correlated with $D_i \cdot \eta_j$, the endogenous variable, since a high value of $p_i^{1900} \cdot c_i^{1900}$ implies that in 1900 the state had high overall city population concentrated in a small number of cities.

The first stage estimating equation using the city-level data is:

$$(\eta_j^{IV} \cdot \ln(D_i)) = \zeta \cdot c_i^{1900} \cdot \eta_j^{IV} + \xi \cdot p_i^{1900} \cdot \eta_j^{IV} + \phi \cdot p_i^{1900} \cdot c_i^{1900} \cdot \eta_j^{IV} + \alpha_i + \gamma_j + +X_{ij}'\Gamma + e_{ij} \tag{2.17}$$

and we hypothesize that $\phi > 0$. States (and below, countries) with a high historical urban population concentrated in a small number of cities should—if the logic of the instrument is correct—have higher population-weighted density today.

State-level IV-2SLS estimates of Equation (2.15) with this second instrumentation strategy are reported in columns 2-3 of Table 2.4.3. The sample is reduced to 39 states because 11 states have no cities in the Chandler (1987) data in 1900. Nevertheless, the estimates remain positive and highly significant. Since p_i^{1900} (total urban population in 1900), one of the excluded instruments, will likely be mechanically correlated with modern popula-

[36] 1900 was chosen because it is the oldest year with broad and global coverage.

Table 2.4.3: State-Level Trade, IV Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Variable is Total Exports from the State-Sector					
Strategy for estimation of density affinity:	η_j computed using industry-level employment			η_j computed using industry-level number of establishments		
D_c x η_j	0.231**	0.149**	0.288***	1.098***	0.657*	0.951***
	(0.0878)	(0.0692)	(0.0865)	(0.408)	(0.381)	(0.349)
ln(population) x η_j			-0.106			-0.0921
			(0.0816)			(0.0738)
K-P F-Statistic	25.159	45.755	25.411	25.251	45.127	37.259
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
States	48	39	39	48	39	39
Observations	4,182	4,132	4,182	4,132	4,132	4,182

Notes: The unit of observation is a state-by-sector pair. The coefficient of interest is the coefficient on an interaction between state-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment in columns 1-3 and estalishments in columns 4-6. All estimates report IV-2SLS estimates. In columns 1 and 3, the excluded instrument is an interaction between sector-level density affinity and state-level population weighted density computed from the US 1900 poplulation distribution. In columns 2-3 and 5-6, the excluded instruments are the total urban population in the state in 1900, the inverse number of cities, and the interaction between the two. The Kleibergen-Paap F-statistic for each first stage regression is reported at the bottom of each column. Standard errors, clustered at the state level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

tion, we control for modern (log of) country population interacted with η_j in column 3; the coefficient of interest remains positive and significant. Finally, in columns 4-6, we repeat the results from columns 1-3 except in all cases use the version of η_j^{IV} estimated from data on establishments rather than data on employment; the results are very similar. Next, we turn to IV-2SLS estimates of our country-level results. Across countries, we rely exclusively on the instruments constructed from the Chandler (1987) city-level data. Although this is a limitation, it is worth noting that across US states, our instrument constructed from the Chandler (1987) data and our direct estimate of historical population weighted density are highly positively correlated; the binned partial correlation plot is reported in Figure 2.A.1.

Country-level IV-2SLS estimates of Equation (2.15) are presented in Panel A of Table 2.4.4; the first stage estimating equation is Equation 3.3.4.1 and first stage estimates are reported in Panels B. For comparison, Panel C reports OLS estimates. Our baseline country-level IV-

Table 2.4.4: Country-Level Trade, IV Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent Variable is Total Exports from the State-Sector					
	Full Sample		Excluding Bottom 10% by Population		Excluding Bottom 10% by Income	
	Panel A: IV-2SLS Estimates					
D_i x η_j	0.517**	0.279**	0.411**	0.319***	0.404**	0.214**
	(0.236)	(0.117)	(0.196)	(0.116)	(0.185)	(0.0894)
ln(population) x η_j		-0.0895**		-0.0434		-0.0887**
		(0.0366)		(0.0407)		(0.0346)
	Panel B: First Stage Estimates					
(p_i, 1900) x (c_i, 1900) x η_j	0.787**	1.021***	0.797**	1.091***	1.119***	1.153***
	(0.344)	(0.312)	(0.338)	(0.345)	(0.382)	(0.356)
p_i, 1900 x η_j	-0.614***	-0.728***	-0.634***	-0.766***	-0.705***	-0.787***
	(0.213)	(0.180)	(0.211)	(0.189)	(0.227)	(0.192)
c_i, 1900 x η_j	-8.705**	-10.54***	-8.782**	-11.36***	-12.43***	-11.83***
	(3.868)	(3.497)	(3.807)	(3.868)	(4.304)	(3.998)
R-Squared	0.095	0.463	0.115	0.474	0.127	0.527
K-P F-Statistic	8.533	27.145	9.104	24.904	9.176	28.569
	Panel C: OLS Estimates					
D_i x η_j	0.134**	0.196***	0.169***	0.181***	0.129**	0.198***
	(0.0624)	(0.0709)	(0.0608)	(0.0676)	(0.0635)	(0.0719)
ln(population) x η_j		-0.0753**		-0.0175		-0.0861**
		(0.0334)		(0.0344)		(0.0332)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Countries	86	86	77	77	78	78
Observations	7022	7022	6281	6281	6379	6379

Notes: The unit of observation is a country-by-year pair. Panel A reports IV-2SLS estimates, Panel B reports first stage estimates, and Panel C reports OLS estimates. The coefficient of interest is the coefficient on an interaction between country-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment. p is the log of the total urban population in 1900 and c is the inverse number of cities. All specifications include country and sector fixed effects, along with other controls listed at the bottom of each column. Sample restrictions are noted in the column header. The Kleibergen-Paap F-statistic for each first stage regression is reported at the bottom of Panel B. Standard errors clustered at the country level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

2SLS estimate is reported in column 1 of Table 2.4.4. The coefficient estimate is positive and significant, supporting the argument that density is a source of comparative advantage and that our baseline estimates are not driven by reverse causality. Reassuringly, and following the state-level analysis, in the first stage specification we find that $\phi > 0$ while the direct effects of p_i^{1900} and c_i^{1900} are both negative. The IV estimate, however, is larger in magnitude than the OLS estimate. One explanation for this is that the IV estimate is capturing a particular local average treatment effect. For example, it could be the case that countries whose modern economic geography is highly correlated with economic geography in 1900 are also countries that industrialized early, and are very specialized in industries that fit their population distribution. This would generate IV estimates that are larger than OLS, since the the IV captures variation across countries whose specialization is most responsive to their population distribution.

Another possible explanation, as noted above, is that variation in the instruments is correlated with the error term in the second stage regression. Indeed, the instruments are constructed from historical population data and likely capture variation in total population and not only variation in D_i. Following the control strategy in our baseline results, in column 2 we include an interaction term between the (log of) present day population and η_j^{IV} as a control. The IV coefficient is smaller in magnitude in column 2 and more precisely estimated. While it remains larger than the OLS estimate, it is no longer statistically distinguishable.

A potential concern with using the Chandler (1987) data is that data quality and coverage are likely different for different sets of countries. In particular, it is likely of lower quality for smaller and lower income countries, which might be more likely to have cities excluded from the data. To make sure this is not driving the result, in columns 3-4 and 5-6 we repeat the specifications from columns 1-2 after dropping countries in the bottom 10% of the population and income distribution respectively. Reassuringly, our estimates remain very similar. The results are also similar if we drop countries in the bottom 20 or 25% of the distribution (not reported).

Taken together, the robustness of our result to the battery of controls and specifications in the previous section, as well as the broadly similar results using these historical instruments, indicates that density is a important and causal determinant of patterns of trade.

2.4.5 Mechanisms: What drives density affinity?

We next turn to potential mechanisms underpinning the baseline results. While in the main specification we relied on a reduced-form measure of industry-level "density affinity," in this section we explore which industry characteristics might drive the baseline estimates. Our approach is to estimate versions of our baseline estimating equation:

$$y_{ij} = \alpha_i + \gamma_j + \xi \cdot \ln(D_i) \cdot Z_j + X'_{ij}\Gamma + e_{ij} \tag{2.18}$$

where Z_j is a vector of sector-level characteristics that potentially determine density affinitty (η_j). We investigate a variety of potential characteristics Z_j. If $\xi = 0$, we interpret that as evidence that Z_j does not drive our main results, whereas if $Z_j > 0$ we interpret that as evidence that Z_j is a potential intervening mechanism. Finally, in order to determine whether our candidate mechanisms can explain our main findings, we add $\eta_j \cdot \ln(D_i)$ to Equation 2.18 and document the extent to which its effect is attenuated by the inclusion of the $Z_j \cdot \ln(D_i)$

First, some recent work has highlighted the greater skill and level of human capital in cities (Davis and Dingel, 2014). In the baseline specification, we control flexibly for the potential role of variation in skill or education, both across sectors and across countries. In column 1 of Table 2.4.5, we report the coefficient on the interaction between population-weighted density and the share of employment in each industry in the US with a college degree. The coefficient on this interaction is positive but statistically insignificant; we also do not find evidence that education is driving the result if we break the industry-level education measure into a larger number of discrete bins (not reported). Another potential determinant of our density affinity measure is the extent to which each sector relies on differentiated local services. Population density might facilitate the productive provision of services and sectors that rely more on local services may therefore benefit disproportionately from density (Abdel-Rahman and Fujita, 1990; Abdel-Rahman and Fujita, 1993). However, we do not find evidence that service reliance explains the export patterns of high-η_j sectors (column 2). The coefficient on the interaction between population-weighted density and industry-level service intensity is in fact negative and far from statistically significant.

Certain industries may locate away from dense cities if they rely on immobile natural resources (e.g. Ades and Glaeser, 1995). These sectors might be less able to benefit from urban externalities and variation in natural resource dependence across industries might drive our variation in density affinity. Indeed, the sectors at the bottom of our "density affinity" list seem to be those that source extensively from natural resources (see Table 2.3.1). To investigate this, we compute the share of natural resource inputs for each manufacturing

Table 2.4.5: Exploring Potential Mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Dependent Variable is Total Exports from the Country-Sector						
D_i x η_j					0.780***	0.281*	0.269*
					(0.183)	(0.143)	(0.163)
D_i x (Share Employment College Educated)$_j$	0.996						
	(1.944)						
D_i x (Services Input Share)$_j$		-0.646					
		(0.620)					
D_i x (Nat. Resource Input Share)$_j$			-1.575**			0.595	
			(0.652)			(0.679)	
D_i x (R&D per Worker)$_j$				0.0844**		0.0242	
				(0.0378)		(0.0276)	
D_i x (Share STEM Workers)$_j$				1.124**		3.938***	
				(0.525)		(0.948)	
Task Controls	No	No	No	No	No	Yes	Yes
Country Level Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry Level Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,437	8,437	8,437	8,333	8,437	8,333	8,437

Notes: The unit of observation is a country-by-sector pair. All specifications include country and sector fixed effects, along with other controls listed at the bottom of each column. Sector-level density affinity computed using the bedrock IV and city-level employment. Additional interactions included in each regression are noted on the left side of the table. Standard errors clustered at the country level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

sector using the US input-output tables. The coefficient on the interaction term between population-weighted density and industry-level natural resource dependence is negative and significant (column 3 of Table 2.4.5), suggesting that indeed denser countries export less in sectors that rely on natural resources. This is consistent with the idea that resource-reliant sectors locate away from urban centers and that dense countries are disproportionately productive in industries that do not rely on natural resources.

Yet another potential mechanism is the role of research and development (R&D) in production. Industries rely differentially on R&D expenditure and innovation in the production process. If cities facilitate innovation (e.g. Duranton and Puga, 2001; Duranton and Puga, 2004), then sectors that rely disproportionately on R&D might be especially productive in dense cities. Our baseline estimates might be capturing the role of density in facilitating R&D. To investigate this, for each sector we compile data on (i) R&D spending per worker and (ii) the share of employees in science, technology, engineering, and mathematical (STEM) fields from the Brookings Advanced Industries database. Again, we include an interaction term between both measures and country-level density in our baseline country-level estimating equation; the estimates are reported in column 4 of Table 2.4.5. Both interactions are positive and statistically significant, suggesting that density may play a role in facilitating R&D and that denser places specialize in the export of R&D intensive sectors.

Finally, we take a more hands-off approach and investigate whether the task content of production in each sector drives the relationship between density affinity and comparative advantage. To measure the task content of each industry, we follow Lanz, Miroudot, and Nordås (2013) and combine data from O*NET on the task content of each occupation with data on occupations by industry from the Occupational Employment Statistics (for the US) and the Labour Force Survey (for the European Union).[37] We aggregate the task content of each occupation to the industry level by weighting each occupation by its share of total employment in the industry (see Section 4 of Lanz, Miroudot, and Nordås, 2013). This yields an industry-level measure of the importance of each of the forty-one O*NET tasks, each of which we interact with country-level density and include on the right hand side of our baseline estimating equation.

While this analysis is necessarily speculative, our main conclusion is that sectors that rely on more interactive and collaborative tasks are disproportionately exported from denser places. The tasks that are important in sectors that disproportionately export from denser countries include "Guiding, Directing, and Motivating Subordinates," "Coaching and Developing Others," "Communicating with Persons Outside Organization" and "Provide Consultation and Advice to Others," and "Selling or Influencing Others." Also in this set are tasks involving technical skill, including "Estimating the Quantifiable Characteristics of Products, Events, or Information" and "Documenting/Recording Information." These findings dovetail with recent work by Michaels, Rauch, and Redding (2019) documenting that since 1880, in the US there has been a dramatic increase in the employment share of "interactive" occupations in metro areas.

Meanwhile, the set of tasks that are significantly less likely to be important in sectors exported by denser countries tend to involve interaction with machines, including "Control-

[37] A potential shortcoming of this approach is the fact that we only have data on the task content of production for the US. Taylor et al. (2008), however, document that the task content of different occupations is very similar across countries.

ling Machines and Processes," "Operating Vehicles, Mechanized Devices, or Equipment," "Repairing and Maintaining Mechanical Equipment." The tasks "Handling and Moving Objects" and "Inspecting Equipment, Structures, or Material" also enter with negative coefficients of similar magnitude; however, they are not statistically significant. The full set of tasks that enter the regression, positively or negatively, with a significant coefficient ($p < 0.1$) are listed in Table 2.A.6.

We next investigate whether these sector-level characteristics drive the effect of density-affinity in our main results. In column 5 of Table 2.4.5 we reproduce the baseline estimate for reference. In column 6, we include controls for all potentially relevant mechanisms described in this section. While the coefficient on the density affinity variable remains positive and (Weakly) significant, its magnitude is reduced by over half, suggesting that the mechanisms described in this section do explain part of the sector-level variation that drives the comparative advantage of denser countries. In column 7, we include only the task content interactions, and the coefficient on the density affinity variable remains similar, suggesting that the task content of more vs. less density-loving sectors form an important underlying mechanism. Nonetheless, it does not fully explain our baseline results, suggesting that additional and un-observed industry characteristics are also at play. Uncovering industry-level characteristics that drive sorting with respect to density strikes us as a potentially interesting area for additional exploration, and we leave a deeper exploration of the determinants of density affinity to future work.

2.5 Conclusion

This paper argues that some countries specialize in density: countries with an abundance of dense cities export relatively more in density-loving sectors Most analysis of sources of comparative advantage in international trade have emphasized aggregate variation in country-level endowments or production technologies. Our theory and empirical results, however, suggest that even when two countries have identical factor endowments in the aggregate, they may specialize in vastly different industries because the domestic distribution of factors of production is a key determinant of comparative advantage. In particular, a key determinant of patterns of trade might lie in the spatial distribution of people within regions and countries.

We first uncover substantial heterogeneity in the density-affinity of tradable sectors, using a new strategy that exploits subterranean geology as a shifter of location-specific population density; while some sectors are disproportionately located in large cities, others are more frequently found in small cities or suburban areas. Next, we show that US states and countries with higher population-weighted density—that is, with a more concentrated population—export relatively more in sectors with high density affinity. Population density and distribution affect not only domestic productivity and inequality, but also comparative advantage and international trade.

The implications of these findings extend into the realms of policy and politics. First, this paper's results suggest that place-based policies might have systematically heterogeneous effects across industries, even to the point of affecting international specialization. By disproportionately benefiting certain places, and perhaps even altering the population distribution, policy could affect sector-level comparative advantage. Second, it is a well-known feature of politics in most countries that more or less dense places achieve different levels of political representation. In the US, for example, institutions like the Senate, the Electoral Collage, and even the lags in House re-districting, lead to the systematic over-representation of less dense areas. Our analysis suggests that this inherently leads to an uneven level of political representation across *sectors*; the resulting political inequality could have major implications for trade policy and the approach to politics that each industry pursues. This interaction between population distributions, political power, and trade is the subject of ongoing work.

Appendices

2.A Figures and Tables for Chapter 2

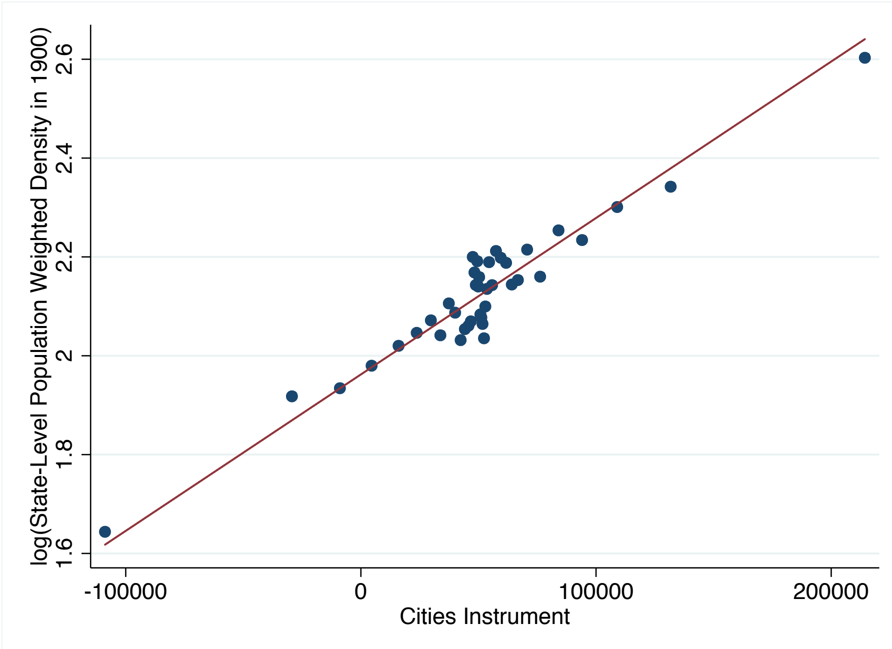

Figure 2.A.1: **Correlation Between Both US State-Level Instruments.**

This figure presents the partial correlation, conditional on state and industry fixed effects, between (i) log of US state-level population weighted density in 1900, estimated from the Fang and Jawitz (2018) data set, and (ii) the interaction between total 1900 city population and the inverse number of cities, estimated from the Chandler (1987) data set.

Table 2.A.1: State-Level Trade, Alternative Specifications

	(1)	(2)
	Dependent Variable is Total Exports (Thousands)	
η_j computed using:		
Employment, IV	0.612***	0.538***
	(0.145)	(0.199)
Establishments, IV	3.508***	3.241***
	(0.541)	(0.660)
Payroll, IV	0.335***	0.295***
	(0.0753)	(0.111)
Employment, OLS	0.788***	0.459***
	(0.236)	(0.172)
Establishments, OLS	2.650***	1.766***
	(0.462)	(0.401)
Payroll, OLS	0.504***	0.307***
	(0.169)	(0.117)
All Controls	No	Yes
State FE	Yes	Yes
Industry FE	Yes	Yes
Observations	4,182	4,132

Notes: The unit of observation is a state-by-sector pair. Each coefficient is an estimate from a separate regression. The coefficient of interest is the coefficient on an interaction between state-level population weighted density and sector-level density affinity using the strategy listed on the left side of the table. All reported specifications are Poisson pseudo-maximum likelihood estimates and include state and sector fixed effects, along with other controls listed at the bottom of each column. Standard errors clustered at the state level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 2.A.2: Main Results: Including Observations with No Exports

	(1)	(2)	(3)	(4)
	US State-Level		Country-Level	
Outcome Variable:	Exports	Exports (asinh)	Exports	Exports (asinh)
Model:	PML	OLS	PML	OLS
D_i x η_j	0.612***	0.425**	0.456***	0.167**
	(0.145)	(0.169)	(0.111)	(0.0720)
State FE	Yes	Yes	-	-
Country FE	-	-	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Observations	4,250	4,250	11,122	11,122
R-squared		0.709		0.823

Notes: The unit of observation is a state-industry pair (columns 1-2) or a country-industry pair (columns 3-4). . The coefficient of interest is the coefficient on an interaction between state- or country-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment. In columns 1 and 3, the outcome variable is total exports and in columns 2 and 4, it is the inverse hyperbolic sine of total exports. Observations with zero exports are included in the estimation. Standard errors clustered at the state (columns 1-2) or country (columns 3-4) level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 2.A.3: Country-Level Trade, Including Continent × Industry Fixed Effects

	(1)	(2)	(3)	(4)	(5)
	Dependent Variable is Total Exports from the Country-Sector				
	Panel A: Outcome Variable is Total Exports (Thousands), PML Model				
D_i x η_j	0.412**	0.403**	0.486***	0.380***	0.491***
	(0.191)	(0.181)	(0.163)	(0.0986)	(0.158)
	Panel B: Outcome Variable is log(Exports), OLS Model				
D_i x η_j	0.139**	0.186**	0.342***	0.179***	0.381***
	(0.0667)	(0.0770)	(0.0757)	(0.0627)	(0.0826)
R-Squared	0.837	0.820	0.821	0.837	0.822
Factor Intensity Controls	No	Yes	No	No	Yes
Country Level Controls	No	No	Yes	No	Yes
Industry Level Controls	No	No	No	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes
Industry x Continent FE	Yes	Yes	Yes	Yes	Yes
Countries	134	90	107	134	83
Observations	10,464	7,159	8,542	10,332	6,674

Notes: The unit of observation is a country-by-sector pair. The coefficient of interest is the coefficient on an interaction between country-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment. Panel A reports Poisson pseudo-maximum likelihood estimates while Panel B reports OLS estimates. All specifications include country and continent-by-sector fixed effects, along with other controls listed at the bottom of each column. Standard errors clustered at the country level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 2.A.4: Main Results: Robustness to Alternative Sources of Population Data

	(1)	(2)	(3)	(4)
Gridded Population Data Set:	LandScan (Baseline)	Global Human Settlement Layer	Gridded Population of the World	Worldpop Project
D_i x η_j	0.456***	0.468***	0.443***	0.497***
	(0.111)	(0.125)	(0.0956)	(0.102)
Country FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Countries	134	134	134	134
Observations	10,547	10,547	10,547	10,547

Notes: The unit of observation is a country-by-sector pair. The coefficient of interest is the coefficient on an interaction between country-level population weighted density and sector-level density affinity computed using the bedrock IV and city-level employment. Population weighted density is computed from a different data set in each column, and the data source is listed at the top of each column. Standard errors clustered at the country level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 2.A.5: Main Results: Alternative Specifications

	(1)	(2)	(3)	(4)
	Dependent Variable is Total Exports (Thousands)			
Sample:	Full Sample		Excluding countries with pop < 1 million	Excluding bottom 10% income
η_j computed using:				
Employment, IV	0.456***	0.774***	0.457***	0.456***
	(0.111)	(0.0720)	(0.111)	(0.111)
Establishments, IV	1.594***	1.836***	1.594***	1.594***
	(0.361)	(0.262)	(0.362)	(0.361)
Payroll, IV	0.248***	0.401***	0.248***	0.248***
	(0.0640)	(0.0408)	(0.0640)	(0.0640)
Employment, OLS	0.292**	0.135	0.292**	0.292**
	(0.147)	(0.0881)	(0.147)	(0.147)
Establishments, OLS	0.793**	0.480**	0.792**	0.792**
	(0.329)	(0.225)	(0.329)	(0.328)
Payroll, OLS	0.224**	0.105*	0.224**	0.224**
	(0.0985)	(0.0580)	(0.0987)	(0.0984)
All Controls	No	Yes	No	No
Country FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Observations	10,464	6,674	9,277	9,515

Notes: All reported coefficients are from regressions at the country-by-sector level. Each coefficient is an estimate from a separate regression. The coefficient of interest is the coefficient on an interaction between country-level population weighted density and sector-level density affinity computed using the strategy listed on the left hand side of each row. All reported specifications are Poisson pseudo-maximum likelihood estimates and include country and sector fixed effects, along with other controls listed at the bottom of each column. Sample restrictions are noted in the column header. Standard errors clustered at the country level, are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 2.A.6: Tasks Associated with Sectors that are the Comparative Advantage of More vs. Less Dense Countries

Panel A: Tasks Associated with CA Sectors in More Dense Countries
Guiding, Directing, and Motivating Subordinates
Coaching and Developing Others
Estimating the Quantifiable Characteristics of Products, Events, or Information
Identifying Objects, Actions, and Events
Selling or Influencing Others
Documenting/Recording Information
Communicating with Persons Outside Organization
Making Decisions and Solving Problems
Provide Consultation and Advice to Others
Panel B: Tasks Associated witrh CA Sectors in Less-Dense Countries
Controlling Machines and Processes
Operating Vehicles, Mechanized Devices, or Equipment
Performing for or Working Directly with the Public
Repairing and Maintaining Mechanical Equipment
Resolving Conflicts and Negotiating with Others
Assisting and Caring for Others
Scheduling Work and Activities
Analyzing Data or Information

Notes: This table lists the set of tasks whose interaction with population weighted density yielded a positive and significant coefficient estimate in Equation 4.4 (Panel A) and the set of tasks whose interaction with population weighted density yielded a negative and significant coefficient estimate in Equation 4.4.

2.B Derivations and proofs for Chapter 2

Housing market

Out of nominal disposable income Y_c, a worker in city c spends a constant share $p_{hc}h_c = \beta Y_c$ on the non-tradable good produced in city c, and a constant share $(1-\beta)Y_c = X_c$ on the basket of tradable sectors, with sub-shares $\alpha_j X_c = p_j c_j^c$ on each sector j. Each landowner faces a price p_{hc} for housing and a cost of P for the numeraire input. Each landowner then uses an amount $X_{hc}(\gamma) = \gamma(1-\xi)(\frac{p_{Hc}}{P})^{\frac{1}{\xi}}$ of tradable inputs, and aggregate housing supply is: $H^s(c) = B_c(\frac{p_{Hc}}{P})^{\frac{1-\xi}{\xi}}$. Equalizing supply and demand yields equilibrium housing prices in each city (equation 2.2):

$$p_{Hc}^{\frac{1}{\xi}} = \beta \frac{L_c Y_c}{B_c P^{\frac{\xi-1}{\xi}}}$$

Landowners in a city receive proceeds from real estate sales $\beta Y_c L_c$, out of which they spend $PX_{hc} = (1-\xi)\beta Y_c L_c$ on the final good, while accruing rents $r_c B_c = \xi\beta Y_c L_c$. r_c is defined as the Ricardian rent per unit of land, increasing in local population density and local disposable income. Using the spatial equilibrium condition and the fact that all land rents are fully rebated to local workers, we have:

$$Y_c = \bar{U} P^{1-\beta} p_{hc}^{\beta} = \bar{U} P^{1-\beta} \left(\frac{\beta L_c Y_c}{B_c P^{\frac{\xi-1}{\xi}}}\right)^{\beta\xi} = \bar{U} P^{1-\beta\xi} \left(\beta \frac{L_c}{B_c} Y_c\right)^{\xi\beta}$$

and thus

$$w_c = P(1-\beta\xi)\bar{U}^{\frac{1}{1-\beta\xi}} \beta^{\frac{\beta\xi}{1-\beta\xi}} \frac{L_c}{B_c}^{\frac{\beta\xi}{1-\beta\xi}} \propto P \times D_c^{\frac{\beta\xi}{1-\beta\xi}}$$

Comparative advantage of cities

Cost minimization by consumers in any location d implies, in the absence of trade costs and using standard Eaton-Kortum algebra (Costinot, Donaldson, and Komunjer, 2011; Michaels, Rauch, and Redding, 2013):

$$p_{dj}(\omega) = \min \{ p_{dcj}(j); c \in C \}$$

The probability that the unit cost is less than p for variety ω of good j produced in c is:

$$F_{jc}(p) = \mathbb{P}(\frac{w_c}{\tilde{z}} < p) = 1 - e^{-(\frac{\frac{w_c}{p}}{A_c D_c^{\tilde{\eta}_j}})^{\theta}}$$

The probability that the minimal cost for variety ω of good j is less than p is thus:

$$F_j(p) = 1 - (\Pi_{c \in C}(1 - F_{jc}(p)) = 1 - e^{-\sum_{c'}(A_{c'} D_{c'}^{\tilde{\eta}_j})^{\theta} w_{c'}^{-\theta} p^{\theta}}$$

and the probability that location c is the lowest cost supplier for variety ω for location d is:

$$\mathbb{P}(\frac{w_c}{\tilde{z}_{jc}} \leq \min\left\{p_{dcj}(j); c \in C\right\}) = \frac{A_c D_c^{\tilde{\eta}_j})^{\theta} w_c^{-\theta}}{\sum_{c'}(A_{c'} D_{c'}^{\tilde{\eta}_j})^{\theta} w_{c'}^{-\theta}}$$

From the Fréchet distribution assumption and the Constant Elasticity of Substitution structure on demand allocation within good j, standard algebra then implies that the share of spending on varieties from location c in sector j must be equal across all locations d:[38]

$$\pi_{dcj} = \pi_{cj} = \frac{p_{cj} X_{dcj}}{X_{dj}} = \frac{(A_c D_c^{\tilde{\eta}_j})^{\theta} w_c^{-\theta}}{\sum_{c'}(A_{c'} D_{c'}^{\tilde{\eta}_j})^{\theta} w_{c'}^{-\theta}} \tag{2.B.0.1}$$

where π_{dcj} denotes spending in city d on goods in sector j produced in city c, equation 2.3 in the model.

Proposition 2.2.1

The derivation borrows from the definition of the unified price index in Redding and Weinstein (2020). Using spending shares 2.3, and the definition of the price index 2.5, we obtain:

$$\frac{\pi_{cj}}{\pi_{cj'}} = (\frac{P_j}{P_{j'}})^{\theta} \frac{(A_c D_c^{\tilde{\eta}_j})^{\theta} w_c^{-\theta}}{(A_c D_c^{\tilde{\eta}_{j'}})^{\theta} w_c^{-\theta}}$$

[38]Given the unbounded nature of the Fréchet distribution, the production structure does not lead to the full specialization of cities in the production of some sectors, which would make the exposition more involved by inducing censoring at the bottom of the sector-city employment density, without adding substantial insight in the model, given that we do not attempt a structural estimation of the parameters

Re-expressing and taking logs on both sides:

$$\frac{\log(\frac{P_j}{P_{j'}}) - (\tilde{\eta}_{j'} - \tilde{\eta}_j)\log(D_c)}{\log(\frac{\pi_{cj}}{\pi_{cj'}})} = \frac{1}{\theta}$$

Multiplying both sides by $\pi_{cj} - \pi_{cj'}$, and using that in autarky $\sum_{c \in C} \pi_{cj} = 1$, summing over all cities c and rearranging yields the Sato-Vartia relative price:

$$\sum_{c \in C} \Big(\frac{\pi_{cj} - \pi_{cj'}}{\log(\pi_{cj}) - \log(\pi_{cj'})}\Big) \log(\frac{P_j}{P_{j'}}) = (\tilde{\eta}_{j'} - \tilde{\eta}_j) \sum_{c \in C} \Big(\frac{\pi_{cj} - \pi_{cj'}}{\log(\pi_{cj}) - \log(\pi_{cj'})}\Big) D_c$$

and, rearranging, we obtain the "Sato-Vartia" relative price expression in proposition 2.2.1.

Population density dispersion

Because equilibrium density D_c is increasing in A_c, at the country level, greater dispersion of A_c therefore leads to greater equilibrium D_c dispersion, as workers reallocate from lower to higher-A_c, higher-D_c locations. The population density distribution in an economy with more dispersed A_c is second-order stochastically dominated by the population density distribution in an economy with less dispersed A_c^* (see 2.B).

Formally, suppose there are two countries, H and F, and define $H(d)$ as the share of the total population living in cities with density below d in H, the high-amenity-dispersion economy:

$$H(d) = \frac{\sum_{c \in C} L_c \mathbb{1}(\frac{L_c}{B_c} \leq d)}{\bar{L}}$$

Let $H^*(d)$ be its counterpart in F. Then, for any d, we have:

$$\int_0^d H(s)ds \geq \int_0^d H^*(s)ds$$

For any percentile p, there is a corresponding density threshold $H^-1(p) = d$. Let the Generalized Lorenz Curve (GLC) of population density be the function:

$$GLC(p) = \int_0^p H^{-1}(q)dq, \text{ for } p \in [0,1]$$

Integration by parts yields:

$$GLC(p) \leq GLC^*(p) \forall p$$

The GLC of density in a country with a higher dispersion of population lies strictly below that of a country with a more concentrated distribution of population. Note that we have, by a change of variable:

$$GLC(p) = \frac{\sum_{c \in C} \frac{(L_c)^2}{B_c} \mathbb{1}(H(\frac{L_c}{B_c}) \leq p)}{\bar{L}}$$

Proposition 2.2.2

We assume, as in Ramondo, Rodríguez-Clare, and Saborío-Rodríguez (2016), that iceberg trade costs are nil within a country, and symmetric (at the country-level) across any two locations in two different countries. The proof follows the structure of Ramondo, Rodríguez-Clare, and Saborío-Rodríguez (2016), extended to a case with many sectors.

We obtain a natural extension of equation 2.4 in a world of many countries, namely that for any city c in country i, the wage bill in sector j satisfies:

$$w_c L_{jc} = \alpha_j \sum_n \frac{(A_c D_c^{\tilde{\eta}_j})^\theta w_c^{-\theta} \tau_{in}^{-\theta}}{\sum_s \sum_{c' \in C_s} (A_{c'} D_{c'}^{\tilde{\eta}_j})^\theta w_{c'}^{-\theta} \tau_{sn}^{-\theta}} \sum_{d \in C_n} w_d L_d \tag{2.B.0.2}$$

We rewrite equation (2.B.0.2) as:

$$w_c = ((\frac{A_c D_c^{\tilde{\eta}_j})^\theta}{L_{jc}})^{\frac{1}{1+\theta}} \Delta_{ij} \tag{2.B.0.3}$$

where Δ_{ij} is a country-sector level variable indexing market access in sector j and country i:

$$\Delta_{ij}^{1+\theta} = \alpha_j \sum_n \frac{\tau_{in}^{-\theta}}{\sum_s \sum_{c' \in C_s} (A_{c'} D_{c'}^{\tilde{\eta}_j})^\theta w_{c'}^{-\theta} \tau_{sn}^{-\theta}} \sum_{d \in C_n} w_d L_d \tag{2.B.0.4}$$

We can use the fact that:

$$\sum_{d \in C_n} w_d L_d = \sum_{d \in C_n} \sum_k w_d L_{dk}$$

and equation (2.B.0.2) to re-express Δ_{ij}:

$$\Delta_{ij}^{1+\theta} = \alpha_j \sum_n \frac{\tau_{in}^{-\theta} \sum_{d \in C_n} \sum_k L_{kd} \big((\frac{A_d D_d^{\tilde{\eta}_k})^\theta}{L_{dk}} \big)^{\frac{1}{1+\theta}} \Delta_{nk}}{\sum_s \sum_{c' \in C_s} (A_{c'} D_{c'}^{\tilde{\eta}_j})^\theta w_{c'}^{-\theta} \tau_{sn}^{-\theta}}$$

$$\Delta_{ij}^{1+\theta} = \alpha_j \sum_n \frac{\tau_{in}^{-\theta} \sum_k \Delta_{nk} L_{nk}^{\frac{\theta}{1+\theta}} \sum_{d \in C_n} (A_d D_d^{\tilde{\eta}_k})^{\frac{\theta}{1+\theta}} (\frac{L_{kd}}{L_{nk}})^{\frac{\theta}{1+\theta}}}{\sum_s \tau_{sn}^{-\theta} \Delta_{sj}^{-\theta} L_{js}^{\frac{\theta}{1+\theta}} \sum_{c' \in C_s} (A_{c'} D_{c'}^{\tilde{\eta}_j})^{\frac{\theta^2}{1+\theta}} (\frac{L_{jc'}}{L_{js}})^{\frac{\theta}{1+\theta}}} \tag{2.B.0.5}$$

where $L_{nk} = \sum_{d \in C_n} L_{dk}$. We define the following objects, that depend on the equilibrium distribution of population within a country:

$$T_{ij} = \Big(\sum_{c \in C_i} (A_c D_c^{\tilde{\eta}_j})^{\frac{\theta}{1+\theta}} (\frac{L_{jc}}{L_{ji}})^{\frac{\theta}{1+\theta}} \Big)^{1+\theta} \tag{2.B.0.6}$$

$$M_i = \sum_j \Delta_{ij} L_{ij}^{\frac{\theta}{1+\theta}} T_{ij}^{\frac{1}{1+\theta}} \tag{2.B.0.7}$$

Note then that we can re-express equation (2.B.0.5) as a system of equations in M_n, T_{sj}, Lsj, and Δ_{sj}:

$$\Delta_{ij}^{1+\theta} = \alpha_j \frac{\sum_n M_n \tau_{in}^{-\theta}}{\sum_s \tau_{is}^{-\theta} \Delta_{sj}^{-\theta} L_{sj}^{\frac{1}{1+\theta}} T_{sj}^{\frac{1}{1+\theta}}} \tag{2.B.0.8}$$

We make note that M_i corresponds to the total tradable wage bill in a country:

$$\sum_{c \in C_i} w_c L_c = \sum_{c \in C_i} \sum_j w_c L_{cj} = \sum_j \Delta_{ij} L_{ij}^{\frac{\theta}{1+\theta}} T_{ij}^{\frac{1}{1+\theta}} = M_i \tag{2.B.0.9}$$

We now use fact (2.B.0.9) to derive the bilateral export flows from country i to country n in sector j, by using the fact that exports of good j from any city $c \in C_i$ to any city $d \in C_n$ are given by:

$$x_{cdj} = \alpha_j w_d L_d \frac{(A_c D_c^{\tilde{\eta}_j})^\theta w_c^{-\theta} \tau_{in}^{-\theta}}{\sum_s \tau_{sn}^{-\theta} \sum_{c' \in C_s} (A_{c'} D_{c'}^{\tilde{\eta}_j})^\theta w_{c'}^{-\theta}}$$

Summing over cities, using (2.B.0.5), (2.B.0.7) and (2.B.0.6), yields, after rearranging:

$$X_{inj} = \sum_{c \in C_i} \sum_{d \in C_n} x_{cdj} = \alpha_j M_n \tau_{in}^{-\theta} \frac{\Delta_{ij}^{-\theta} T_{ij}^{\frac{1}{1+\theta}} L_{ij}^{\frac{\theta}{1+\theta}}}{\sum_s \Delta_{sj}^{-\theta} T_{sj}^{\frac{1}{1+\theta}} L_{sj}^{\frac{\theta}{1+\theta}}} \tag{2.B.0.10}$$

We next derive the average wage in country i and sector j:

$$w_{ij} = \frac{\sum_{c \in C_i} w_c L_{cj}}{\sum_{c \in C_i} L_{cj}}$$

by using equation (2.B.0.2), again summing over all cities in country i and using the same manipulations:

$$w_{ij} = \frac{\sum_n X_{inj}}{\sum_{c \in C_i} L_{cj}} = \frac{\sum_n X_{inj}}{L_{ij}} = \alpha_j \frac{\sum_n M_n \tau_{in}^{-\theta} \Delta_{ij}^{-\theta} T_{ij}^{\frac{1}{1+\theta}} L_{ij}^{-\frac{1}{1+\theta}}}{\sum_s \Delta_{sj}^{-\theta} T_{sj}^{\frac{1}{1+\theta}} L_{sj}^{\frac{\theta}{1+\theta}}} \tag{2.B.0.11}$$

and, using the system (2.B.0.8) and substituting, we obtain:

$$w_{ij} = \Delta_{ij} (\frac{T_{ij}}{L_{ij}})^{\frac{1}{1+\theta}} \tag{2.B.0.12}$$

Plugging (2.B.0.12) into equation (2.B.0.10) yields proposition 2.2.2.

Chapter 3

Testing the Self-Interested Voter Hypothesis [1]

3.1 Introduction

American journalist and satirist H.L. Mencken once quipped that *"every election is a sort of advance auction sale of stolen goods"*[2]. The cynicism of that proposition notwithstanding, theories of political economy often start from the widely accepted tenet that rational voters elect candidates based on their narrow self-interest, and that expected material gains from one policy platform over another could push marginal voters to choose a side in electoral contests.[3] During local or national campaigns, candidates regularly promise redistributive transfers, reductions in taxes, or "pork-barrel" government spending targeted towards certain locations, production factors, or specific segments of the electorate. Whether these electoral promises mostly constitute "cheap talk" of little credibility, or whether campaign announcements shape actual voting behavior, is a crucial question for the empirical content of economic models of the vote. Can politicians win support by merely promising higher

[1] I am grateful to Matteo Alpino, David Amiel, Alexandre Andorra, Felipe Carozzi, Elizabeth Cascio, Celine Colange, William Fischel, Joel Gombin, Etienne Lehmann, and Francesco Passarelli, as well as participants at the 2020 virtual meetings of the Urban Economics Association for helpful discussions and feedback on this chapter. Brian Liu provided excellent research assistance. Part of this chapter was written while in residency at Dartmouth College and at CEPREMAP - the hospitality of both is gratefully acknowledged.

[2] *Sham Battle*, The Baltimore Evening Sun, October 26, 1936.

[3] See e.g. the seminal contributions of Downs (1957), Lindbeck and Weibull (1993), and Lizzeri and Persico (2001); or Persson and Tabellini (2016) for a general treatment of economic voting models.

transfers or lower taxes on certain constituencies once elected?

In this paper, I leverage quasi-random variation in exposure to a local tax cut across more than 35 000 French townships (*communes*) to assess the causal impact of a campaign promise on voter support. In the months leading to the highly disputed first round of the 2017 French presidential election, newcomer candidate Emmanuel Macron promised to repeal a widely unpopular housing levy, the "*taxe d'habitation*" (TH), worth close to 2 percent of disposable household income on average, while pledging the central government would compensate municipalities for any lost revenue. The promised repeal of the TH, combined with the implicit assurance that no additional local taxes would need to be raised to replace foregone receipts, created considerable spatial heterogeneity in expected net savings from this nationwide policy platform, depending on the initial TH tax bill. Using administrative data on variation in the local TH burden, and polling-place-level electoral returns, I show that expected benefits from the tax rebate causally explain a substantial share of Macron's electoral support, over and beyond his predicted success based on the characteristics of the local voting body.

The pre-reform housing tax bill in a French town may plausibly have been correlated with observed and unobserved determinants of the propensity of the municipality's electorate to lean towards Emmanuel Macron. For example, due to spatial sorting by political preferences across towns with varying housing values, potential Macron voters might merely be more likely to locate in high-price locations.[4] Alternatively, city councils in towns with more Macron supporters may have been more prone to set high initial housing tax rates, again correlating the initial tax burden with local Macron support, independently of the effect of the promise itself. As described in detail in section 3.2, however, the tax base for the TH was devised in reference not to current housing market values, but to outdated appraisals assessed in the early 1970s for a few "reference units" in each town, never updated since besides a nationwide inflation adjustment, and outside the control of today's local authorities. These institutional features implied large variation in initial tax burdens between municipalities otherwise comparable in current days, depending on idiosyncratic 1970 assessments. "Assessed rental values" from the 1970s, averaged at the municipal level, thus

[4]See e.g. Rodden (2010) for a recent overview of this phenomenon and its electoral consequences in the United States.

constitute an instrument for the expected local tax savings, both strongly relevant – as they directly enter the tax computation, and plausibly orthogonal to other local determinants of the propensity to vote for Macron, conditional on current housing values interacted with current local tax rates.

Adjusting for housing market values – computed from exhaustive transaction data on millions of pre-election home sales – and administrative measures of local tax rates, I show that voters were more likely to choose Emmanuel Macron when they expected his election to offset a higher initial local tax bill. Higher expected savings from the TH repeal, driven by variation in 1970 assessments of the tax base, causally increase the local Macron vote share. The magnitude of the effect is large: municipalities with a one standard deviation higher tax bill (c. 1 percent of median household income) exhibit a statistically significant improvement of Emmanuel Macron's score of 1 to 1.5 percentage point, conditional on his predicted share of the vote based on electoral results from past presidential elections and controlling for socio-demographic characteristics.

Heterogeneity analysis using Census data lends credibility to this causal mechanism. Localities with a larger share of homeowners (who stood to gain more, through capital gains, from reduced user costs of housing), and those with a smaller share of households exempt from the tax, were more responsive to the policy promise. The results survive a battery of robustness checks, ranging from additional observable controls, to comparing only municipalities that share a geographic border in order to adjust for spatially correlated political shifts, to using other electoral contests as placebo tests. Moreover, I show that voter search for information about Macron's proposal on the housing tax, as inferred from Google searches, spiked around the announcement and close to election day. Event studies, using both aggregate high-frequency polling and prediction markets data, also display a nationwide impact of the tax cut announcement on Macron's winning chances, consistent with the aggregation of estimated local effects.

Therefore, using the conditionally random component of pre-treatment exposure to the planned cut in the housing tax at the municipality level, across more than 35,000 towns, I validate the hypothesis that material self-interest constitutes a key motivation for voting behavior, even in ideologically polarized contests such as the 2017 French presidential election, hailed worldwide as a defining showdown between a populist far-right animated by identity concerns and a progressive center.[5] The results suggest that tax cuts aimed at peculiarly salient, lump-sum, and politically controversial levies, might be especially efficient in triggering increased electoral support.

Politicians regularly make campaign promises involving substantial amounts of fiscal and non-fiscal redistribution across income levels, age groups, or other constituencies. A considerable body of theoretical work has studied the pre-election role of such "distributive politics" (Weingast, Shepsle, and Johnsen, 1981), and the resulting incentives for politicians to engage in material campaign promises (Feldman, 1982; Aragonès, Postlewaite, and Palfrey, 2007). While it is commonly accepted that voters may make electoral decisions based on their own self-interest, credible empirical evidence is surprisingly scant, and this paper contributes to our understanding of this phenomenon.

On the one hand, a large empirical corpus, summarized early on by Lewis-Beck and Stegmaier (2000), has explored "retrospective" economic determinants of the vote. Existing work has documented the role of cyclical economic outcomes, or windfall gains, for the electoral success of incumbents, at the national or local level, using time series evidence (Brender and Drazen, 2008), or, more recently, quasi-random spatial shocks to income (Brunner, Ross, and Washington, 2011) or wealth (Bagues and Esteve-Volart, 2016) to test for the role of "attributional" motives in the incumbent vote share. Other studies have examined the specific role of past fiscal transfers for electoral outcomes, using quasi-experimental variation in expenditure outlays (Levitt and Snyder Jr, 1997; Huet-Vaughn, 2019), or truly randomized cash payouts (Manacorda, Miguel, and Vigorito, 2011; De La O, 2013; Zucco Jr, 2013), but tax cuts have received comparatively less attention. A handful of contributions have exploited local variation in exposure to a nationwide policy, with the goal of estimating their *ex post* electoral impact for local incumbents, a distinct question from my focus on campaign promises and their *ex ante* impact. Most recently, Casaburi and Troiano (2016) show that an increase in tax enforcement intensity in Italy, hitting various regions with differential intensity, increased local reelection probabilities; and Fetzer (2019) finds that higher exposure to fiscal austerity in the United Kingdom increased local support for the UKIP party and Britain's exit from the European Union. This strand of literature has heretofore focused on "retrospective" motives, whereby voters reward or punish incumbents based on past transfers or taxes. While

[5] See e.g. *Financial Times*, April 26, 2017, "Emmanuel Macron offers the patriotic antidote to nationalism".

I also rely on plausibly exogenous spatial variation in policy exposure, my focus is on the *ex ante* response of the vote to a campaign promise in a national - rather than local - electoral contest, before the program is implemented.

Only a few studies have explored the effect of campaign promises of *future* tax and transfers on electoral support. Recent work on campaign efficacy has taken advantage of differential quasi-random or truly randomized exposure to *information* about politicians' platforms (Kendall, Nannicini, and Trebbi, 2015; Cruz et al., 2018; Spenkuch and Toniatti, 2018), focusing on the role of voter knowledge but leaving aside differential exposure to actual policies implemented. Alpino (2018) studies a 2006 Italian electoral promise of tax cuts for homeowners, and finds a positive impact on voting intentions for Silvio Berlusconi in surveyed homeowners and in localities with higher home-ownership. Work by Elinder, Jordahl, and Poutvaara (2015) finds in declarative survey data that parents of younger children were more likely than their less affected counterparts to vote for the Social-Democrat Party when it promised them increased transfers. My strategy exploits an instrumental variable approach, combined with quasi-random geographic variation in pre-policy exposure to a promised tax cut; it relies on exhaustive town-level electoral returns and administrative tax data, and is concerned with a large tax cut proposal representing close to 2 percent of household annual income. This paper thus expands our understanding of the electoral efficacy of campaign promises at the national level, adjusting for predictable ideological shifts and socio-demographic composition, and leveraging plausibly exogenous local variation.[6]

3.2 Institutional background and data

3.2.1 The 2017 French presidential election

I study the electoral benefits of material campaign promises in the context of the 2017 presidential election in France. France has a democratic presidential regime; the president is elected every five years by popular vote in a two-rounds majoritarian electoral system. In

[6]By focusing on the differential electoral impact of a local residence tax repeal, this paper also speaks to the literature on the political economy of local taxation. Existing work suggests that municipalities reduce local tax rates in election years. While several authors have investigated these local political budget cycles (e.g. Aidt and Mooney (2014) in London, Foremny and Riedel (2014) in Germany, or Alesina and Paradisi (2017) in Italy), the role of local taxes for national electoral outcomes remains mostly *terra incognita*.

April 2017, the first round pitted eleven candidates against each other. The 2017 election exhibited a series of unusual features, making it especially well-suited for the study of the electoral efficacy of campaign promises.

First, incumbent Socialist president François Hollande, facing historically low popularity ratings, decided not to run for re-election on December 1st, 2016. A rare occurrence in French politics, the decision is helpful to my empirical design, as the absence of an incumbent candidate voids the usual difficulty associated with disentangling retrospective voting from the impact of campaign promises. Second, two major traditional parties, the center-right *Les Republicains* (*LR*) and left-leaning *Parti Socialiste* (*PS*), held open primaries for the first time since the beginning of the Fifth Republic in 1958. Both chose "surprise" candidates, former Prime Minister François Fillon for *LR* and former Education Minister Benoit Hamon for *PS*, instead of the early poll leaders, former Prime Ministers Alain Juppe (*LR*) and Manuel Valls (*PS*). Third, François Fillon, the candidate for the main opposition party (*LR*), was engulfed in a political-financial scandal.[7] News coverage of the misconduct made Fillon, the early favourite in January, drop to third place in a few weeks, and generated substantial volatility in polling intentions and prediction markets.

In this context of heightened political uncertainty among traditional parties, Emmanuel Macron, a former Economy Minister in Hollande's cabinet, but a newcomer in electoral politics, mounted a bid for the presidency. He founded a new political party, *En Marche*, in April 2016, kick-starting his campaign by a show of force in his hometown, Amiens. He formally announced his candidacy on November 16th, 2016. Fillon and Hamon had run primary campaigns to the right- and left-wing of their respective parties, opening space for a centrist candidacy, according to observers[8]. On February 24th, 2017, Emmanuel Macron was a guest of France's most watched morning show, *Bourdin Direct*. He announced that, if

[7] *Le Canard Enchaine*, a satirical weekly newspaper, revealed in January 2017 that Mr Fillon had employed his wife and two of his children as aides while a member of parliament, and that they were paid in exchange for little actual work.

[8] Daniel Boy, a researcher at the CEVIPOF, France's largest political science research center, remarked in January 2017: "*Between these two candidates, substantial political space is emerging in the center and the context is ideal for Emmanuel Macron.*" Historian Jean Garrigues concurred: "*Part of voters who supported* [centrist candidates] *Alain Juppe in the LR primary and Manuel Valls in the PS primary feel incompatible with the primary winners. There is a programmatic and political space in between.*" *Présidentielle: Entre Fillon et Hamon, Macron a-t-il un boulevard au centre?*, 20 Minutes, January 31st, 2017. Competition in the general election was thus open, with five candidates (Fillon, Hamon, Macron, Marine Le Pen, leader of the far-right *Front National-FN*, and Jean-Luc Melenchon, chief of far-left party *La France Insoumise-LFI*) enjoying plausible chances of advancing to the runoff stage.

elected president, he would exempt the bottom 80% of households by income (c. 24 million households) from the *taxe d'habitation*, a local residence tax. On April 23$^{\text{rd}}$, Macron received the most votes in the first round of the presidential election, with 24.01% of the overall vote, advancing to the runoff stage where he faced populist right-wing candidate Marine Le Pen. On May 7$^{\text{th}}$ 2017, Macron won the second round of the presidential election with c. 66.1% of the vote, becoming the eighth elected president of the Fifth Republic.

3.2.2 The French housing tax

The *taxe d'habitation* (TH) Mr Macron promised to repeal is a housing tax assessed by the State on behalf of local governments. It applies to all furnished housing units and is remitted by resident households, irrespective of their ownership status (tenants or owner-occupiers). It is one of four main sources of recurring tax revenue for local authorities in France.[9] While the TH follows a complex array of rules, it is broadly understood as the product of historical rental values determined in 1970 by the sum of two tax rates set at the municipal level. I describe its main features of interest for the empirical strategy succinctly below, and provide additional details on its computation in appendix 3.D.

Tax base The tax base is an assessed rental value of the unit, the *valeur locative cadastrale* in year n (VLCn). The VLC is not predicated on the market value of the unit. Rather, it is the product of an estimated *weighted square footage* (*surface ponderee*), and a *reference rent* per square meter (*tarif de reference*):

- The *weighted square footage* re-weights the actual floor area of a unit in 1970 (or at the time of its construction) according to a complex parametric formula depending on a quadratic polynomial of area, eight "quality ladders", original 1970 condition and neighbourhood status, and the presence of what were deemed in 1970 to be "comfort" elements, such as running water or an elevator.
- The *reference rent* is based on the actual rent (per square meter of weighted square footage) of a few "representative units" (5.2 percent of all units nationally), recorded as

[9] These also include a property tax (the *taxe fonciere sur les proprietes baties*), a tax on undeveloped land (*taxe fonciere sur les proprietes non-baties*), and a local business tax (the Local Economic Contribution or *Contribution Economique Territoriale*).

of January 1st, 1970, for each quality ladder in most towns. All other dwellings were allocated the corresponding municipality-quality-cell specific reference rent.

Multiplying the "reference rent" by the "weighted square footage" of the unit yielded a nominal assessment (VLC) as of January 1970, known as the *VLC70*, for each unit. In theory, the law calls for a revision of the *VLC70* every three years. In practice, revisions never occurred, due to the administrative complexity of the task.[10] VLCs were "actualized" once in 1980 at the province-level, yielding *VLC80*.[11]. Since 1981, VLCs are increased by an annual indexation coefficient applicable *nationwide*, yielding a year *t VLCt* for each unit. Special deductions off that tax base are available for older citizens, lower-income households, people with disabilities, and households with children.[12] Formally, the tax base of a housing unit in quality ladder k and municipality j, occupied by household h with characteristics h in year t is:

$$\text{Base}_{hjkt} = \underbrace{\overbrace{\text{WeightedSqFt}_{hk} \times \text{Rate}_{jk,1970}}^{\text{VLC70}_{jh}} \times \text{A}_{\text{Dep}(j),70\to 80} \times \text{I}_{1980\to t}}_{\text{VLC}_{jht}} - \underbrace{\sum_r \text{DedRate}_{rjt}(X_{ht}) \times \overline{\text{VLC}}_{jt}}_{\text{Deductions}_{hjt}} \tag{3.2.2.1}$$

Tax rates Municipalities receive information from the national revenue service only on the aggregate VLC in the town, not the VLC of individual units. Given their absence of control over the tax base, the rate of the TH was one of the instruments available to city councils to balance their budget until 2017. As of 2016, there were two main rates in force, at the municipality level, and at the level of the inter-municipal cooperation (a syndicate of towns jointly providing local public services, known as *Etablissements publics de coopération intercommunale*) – with the larger share attributable to the municipality. Tax rates vary widely across towns, reflecting political priorities and funding needs, conditional on aggregate assessed values. The distribution of tax rates as of 2016 is described in figure 3.A.1, and exhibits substantial

[10]In rare cases, changes to the imputed square footage of a unit were made in the case of remodellings or additions.

[11]There are 96 provinces or *departements*, each of which contains c. 350 municipalities on average, making the 1980 adjustment imperfect and noisy, and not unit-specific. More details are provided in appendix 3.D

[12]They take the form of a fixed rate (e.g. 10% per additional child) multiplied by the *average* VLC in the municipality, $\overline{\text{VLC}}_{jt}$.

dispersion across municipalities.

Summary of the tax computation The total tax due by a household is increased by a flat fee F_t of 1 to 3 percent of the tax due, to fund the administrative collection workload for the central government. Additional levies accruing to the central government are applied to "luxury" units with a VLC above some threshold $\bar{B}$. Some tax filers (people above 60, disabled individuals, and widowers) with an income below a threshold close to the national poverty line adjusted for household size, are exempt from the tax. Finally, the tax is subject to a ceiling of 3.44% of annual taxable income. A household h with income Y_{ht} and characteristics X_{ht}, residing in municipality j in year t, in a unit classified in category k, thus owes:

$$\text{Tax}_{hjkt} = \begin{cases} 0 & \text{if } \text{Y}_{ht} \in \Omega_t(X_{ht}), \text{ the exoneration set for households with characteristics } X_{ht} \\ \min\left\{[1+\text{F}_t] \times \left(\tau_{jt}^{\text{Com}} + \tau_t^{\text{EPCI}(j)} + \tau_t^{\text{High}} \mathbb{1}[\text{Base}_{hjkt} \geq \bar{B}]\right) \times \text{Base}_{hjkt}\, ; 0.0344 \times \text{Y}_{ht}\right\} & \text{otherwise} \end{cases} \tag{3.2.2.2}$$

where the base is Base_{hjkt}, computed in equation 3.2.2.1.

Heterogeneity of the tax burden The absence of revisions of appraised values over more than 45 years led to a substantial disconnect between market valuations and VLC assessments. Housing units in some regions, such as the Mediterranean and Atlantic coasts, and homes in suburban areas, were newly built in 1970. They benefited from then "modern" amenities, leading to high reference rents in what had become mostly decayed social housing units or periphery areas forty-seven years later, in 2017. On the contrary, city centers in most of France's large cities were only rehabilitated in the 1980s and 1990s. Along with newly gentrified rural towns, these areas had low *VLC70*s, both due to their poor condition in 1970, and to their relatively unattractive geographic situation back then. The TH was thus widely decried as France's most unfair tax, mostly due to horizontal equity concerns.[13] The tax burden was also substantially lesser in recently renovated apartment buildings in city-centers, mostly populated by higher-income households, than in single-family homes in rural periphery towns and in large social housing projects in urban suburbs. This burden-

[13] See e.g., among countless examples, *Challenges*, a popular weekly, on November, 13th, 2014, titling: "Why the housing tax is the most unfair of all taxes".

shifting regularly led to Parliamentary reports, the creation of advisory committees, and popular press outcry. The tax was especially burdensome for younger and lower to middle-income households, given its applicability to tenants, and its one-time annual lump-sum payment.

3.2.3 Data description

I describe succinctly below the use of electoral results, local government finance reports, administrative house sales data, Census counts, and spatial geo-coded information for the construction of my main sample. Additional details on the construction of the data are provided in appendix 3.C.

To construct my outcome variables, I use exhaustive data on electoral outcomes for all municipalities in France in the first rounds of the 2017, 2012, and 2007 presidential elections from the French Interior Ministry. I use nationwide polling data from *Elabe* to predict, using a shift-share strategy, local Macron support from aggregate polling shifts from 2012 candidates to 2017, as explained in section 3.3.

To construct my measure of exposure to the tax rebate, I use the French Treasury's report on local taxes as of 2016, the last year of taxes known to households at the time of the election. I obtain TH rates, revenues, aggregate and average rental values (VLC) for each town, as well as the municipality-specific amount and number of beneficiaries of all deductions and exemptions. Figure 3.A.3a displays spatial variation in my baseline exposure measure (average TH receipts per household subject to the tax), across 35,197 municipalities merged to electoral data which constitute the baseline sample.[14] Across these main-sample townships or *communes*, with a median number of registered voters of 345 (mean: 1,260), the average housing tax burden (weighted by registered voters) was c. 733 euro, with a standard deviation of 245 euro. On average, it represented 2.3 percent of household income.

When using historical assessments as an instrument for the local tax burden, I must control for current house prices, in order to adjust for household spatial sorting. I construct city-level mean values from exhaustive transaction data on c. 5 million home sales in a five-

[14]These 35,197 municipalities on French mainland territory represent 99.97 percent of towns and close to 94 percent of the electorate in the 2017 election. When including additional controls, the sample size is reduced due to statistical secrecy procedures.

year period centered in 2017.[15] Current house prices are correlated with the mean VLC in a town, as shown in figure 3.A.4: long-run persistence in amenities and productivity implies auto-correlation in home values, even over close to five decades (Rosenthal and Ross, 2015). Nonetheless, there remains substantial heterogeneity in VLCs, conditional on current home values.[16] I control for local socio-demographic characteristics from a variety of sources: deciles of household income, capital, pensions, and labor shares of income, aggregate tax payments, the share of homeowners and the share of secondary residences in a municipality. To compare only contiguous municipalities, I use adjacency lists from *OpenStreetMap*. Finally, to study the dynamics of electoral odds around the TH repeal announcement in February 2017, I employ information from a French polls aggregator website, *PollsPosition*, Google Trends Internet search data, and prediction market prices from PredictIt.
Table 3.B.1 describes summary statistics of my sample across municipalities for two subsets of controls: those used in the full sample (panel (a)), and those restricted to the subset of larger municipalities (or county seats, *chefs-lieux de canton*) with available details on the full distribution of household income by decile (panel (b)).

3.3 Conceptual framework and estimation strategy

3.3.1 Causal inference on voter choice

Measuring the efficacy of campaign promises in delivering electoral support is a crucial step in giving empirical credence to the theory of the self-interested voter, but identification issues can plague its direct estimation. Let PP_G^C be the monetary benefit that group G draws from a policy promise by candidate C, and V_i^c the event where voter i votes for candidate C. A simple model of voter choice whereby voters vote sincerely to maximize utility, subject to

[15]Because of different local laws dating back to the German occupation of France in the 19th century, the DVF database does not cover c. 3000 municipalities in three provinces in the easternmost part of France, which are therefore excluded from the sample when using the instrumental variable approach.

[16]A regression at the municipality level of log VLC on log median home value yields an R-squared of only 0.32 (0.27 when using the mean home value), suggesting substantial unexplained variance in mean appraised values across towns conditional on current real estate market conditions.

some idiosyncratic noise, suggests estimating an individual-level regression:

$$\mathbb{P}(V_i^c = 1) = \alpha + \beta PP_G^C \times \mathbb{1}(i \in G) + \epsilon_i^C \tag{3.3.1.1}$$

Conceptually, the parameter of interest is β: the effect of an additional euro promised by candidate C to group G on the probability that a voter $i \in G$ votes for C. However, the benefit of voting for C, from the point of view of voter i, is not randomly assigned. Merely quantifying whether potential beneficiaries of a campaign promise are more likely to support the corresponding candidate is subject to a fundamental endogeneity concern: ϵ_i^C can be correlated with $PP_G^C \times \mathbb{1}(i \in G)$. This issue can be traced to five main sources.

Measurement error First, neither platforms, nor policy preferences, are uni-dimensional. The determinants of each constituency's voting decision, whether economic or not, are manyfold, and identifying the precise impact of one electoral promise ($PP_G^C \mathbb{1}(i \in G)$), out of the many that constitute a politician's platform ($\sum_{G'} PP_{G'}^C \mathbb{1}(i \in G')$, is a difficult task, when there is correlation both across policy promises ($Cov(PP_{G'}^C, PP_{G'}^C \neq 0$) and in membership of various constituencies ($Cov(\mathbb{1}(i \in G'), \mathbb{1}(i \in G)) \neq 0$). Conversely, because policies have complex consequences for individual situations, it is often hard to pinpoint *ex ante* precisely who stands to gain from a given material campaign promise, and thus to link targeted groups to their electoral response. This latter difficulty is especially pronounced when using survey data, where not only voter choice is declarative and noisy, but only limited demographic information (whether $\mathbb{1}(i \in G) = 1$) is available to quantify respondents' potential gains from a platform. Both group membership, and benefits from a candidate's platform, are thus *measured with substantial error*, attenuating any directional relationship between political support and potential financial rewards.

Reverse causality Second, policy promises are not randomly assigned, but rather the outcome of a strategic decision by candidates, which can reflect expected electoral strength. A candidate could try to mobilize potential supporters, or, conversely, to counter weakness in a given category of voters, by offering a policy PP_G^C designed to benefit them. Thus, *reverse causality* may amplify or attenuate the correlation between individual benefits from a policy platform, on the one hand, and the likelihood of voting for a candidate, on the other.

Confounding Third, voter behavior could be partly determined by an omitted variable, a component of ϵ_i^C, correlated with policy platforms, confounding the direct causal relationship. Non-economic rationales for a specific policy may preclude a direct estimation of the impact of monetary gains from a platform on electoral support. For instance, a candidate may promise to impose international tariffs for "identity"-related reasons (Grossman and Helpman, n.d.); if poorer individuals consume more imported goods (Fajgelbaum and Khandelwal, 2016), but are also more likely to vote based on identity motivations (Rodrik, 2020), they could be found to support the tariff candidate, in spite of their purchasing power being harmed more by her policies. Empirically, this means that unobserved determinants of voting behavior, such as personal ideology, tend to be correlated with membership in groups targeted by specific policy platforms, leading to *omitted variable bias*.

External validity Fourth, and partly due to the aforementioned issues, researchers wishing to identify the causal impact of electoral promises sometimes resort to location-specific campaign issues (e.g. targeted spending on local infrastructure, such as a new airport in Ahlfeldt and Maennig (2015)), where measurement error or reverse causality are a lesser concern. However, an trade-off arises between external validity and bias: using well-identified local variation in electoral promises may preclude drawing broader implications for nationwide tax and spending policy platforms. Indeed, because locally targeted spending, by its very nature, rarely represents a substantial share of a country's aggregate income, generalizing its consequences to electoral responsiveness to national policies is challenging.

Retrospective and prospective motives Finally, the presence of *retrospective* voting motives ("punishment" and "rewards" for past behavior in office) makes the role of incumbent candidates challenging. Officeholders may pledge distinct policies, and also stand different chances in electoral contests, due to incumbency advantage or fatigue. Incumbency can then hamper the estimation of the role of policy promises in triggering electoral support: voter motivations may include both retrospective and prospective dimensions; and the credibility of pledges may vary between incumbents and other candidates.

3.3.2 Baseline specification

Overcoming these limitations thus requires precisely measured, economically substantial, and one-dimensional variation in the benefits expected by different groups from a nationwide policy promise. We need credibly exogenous variation in gains across constituencies to be neither confounded by correlates of ideology, nor driven by political targeting, and operate in the absence of any incumbency effect. Spatial variation in pre-reform exposure to the housing tax repeal in the 2017 French presidential election leverages exactly such heterogeneity.

Defining $j(i)$ as the municipality where voter i lives, a promise by candidate C to repeal the local housing tax effectively amounts to a locally-varying policy: $PP_j^C \times \mathbb{1}(i \in j) = TH_{j(i)}$, so that equation 3.3.1.1 writes (adding individual observed determinants of the vote X_i) :

$$\mathbb{P}(V_i^c = 1) = \alpha + \beta TH_{j(i)} + \gamma X_i + \epsilon_i^C \tag{3.3.2.1}$$

Aggregating votes at the municipal level yields a relationship between the share of votes $\omega_{j,2017}$ for candidate C (Macron) and the local TH level TH_j. Under the self-interested voter hypothesis, voters located in towns facing an initially higher housing tax burden, should vote in higher numbers for Emmanuel Macron, motivated by their narrow material interests, all else equal:

$$\omega_{j,2017} = \alpha + \beta TH_j + \gamma \bar{X}_j + \bar{\epsilon}_j^C \tag{3.3.2.2}$$

By using administrative data at the town level, and exploiting pre-determined exposure to the nationwide promised tax repeal in the absence of incumbency, the specification in equation 3.3.2.2 addresses the measurement error, external validity, and retrospective voting issues exposed above. However, the confounding concern remains: unobserved determinants of local political ideology might be correlated with the pre-reform housing tax exposure in a municipality. If locations with higher tax burdens also have electorates with different propensities to vote for Macron, independently of the TH repeal promise, omitted variables at the town level may bias estimates of the causal impact of the pledge. To disentangle the specific role of the tax cut proposal, I proceed in two steps. I first adjust for persistent observed and unobserved determinants of voter behavior; I then use a plausibly exogenous driver of the TH burden to instrument spatial variation in pre-reform exposure.

3.3.3 Adjusting for local ideology

The first approach adjusts the electoral outcome measure (the Macron vote share in the first round of the 2017 election) for two alternative proxies of the predicted local Macron vote share, the "polling shift" and the "prediction error" methods. Both methodologies use as their dependent variable, instead of the raw Macron vote share, the deviation of local Macron support from its expected value based on different prediction methods:[17]

- *The polling shift approach*. This methodology predicts the 2017 local Macron vote based on a shift-share design, using aggregate shifts in support from 2012 to 2017 *before* the policy announcement, combined with local 2012 results. A study released on February 21st study[18] included a decomposition of voting intentions by 2012 first-round vote of the respondent (for the top five 2012 contestants). I construct the polling shift residualized variable as the difference between the actual Macron vote share in municipality j, $\omega_{j,Mac,2017}$, and the predicted share based on nationwide shifts from 2012 candidates to Macron in 2017:
$$PS_j = \omega_{j,Mac,2017} - \sum_{k=1}^{5} \lambda_k \omega_{jk,2012}$$
where $\omega_{jk,2012}$ is the share of 2012 registered voters who supported candidate k in municipality j, and λ_k is the share of candidate k voters supporting Macron in 2017. PS_j is the "excess" Macron support, in city j, relative to his predicted strength based on support for any of the top five 2012 candidates *before* he announced the TH exemption. [19]

[17]While Macron was not a candidate in previous elections, making a difference-in-differences strategy infeasible, adjusting his vote share for his predicted support based on past elections has a logic similar to a DiD strategy. Section 3.4.3 explores this further using past elections as placebo tests.

[18]I rely on pre-announcement polling data from Elabe, a leading French pollster. Fieldwork was conducted five to three days prior to the tax cut announcement.

[19]The polling shifts λ_k are: 3 percent of 2012 Melenchon voters, 34 percent of Hollande voters, 42 percent of Bayrou voters, 15 percent of Sarkozy voters, and 1 percent of Le Pen voters, intending to vote for Macron in the first round. Since the top five candidates in 2012 represent 95 percent of the vote nationwide, I re-scale the predicted value so that the aggregate nationwide number of predicted Macron votes in 2017 matches the actual number of Macron voters. The linear combination of the top five candidates in 2012 has high predictive power for the actual 2017 Macron vote, with an R^2 of 36 percent.

- *The prediction error approach*. In this alternative methodology, I regress the 2017 vote share of Emmanuel Macron on the vote share (as a share of registered voters) of each of the 10 candidates (plus abstention) in the first round of the 2012 election.

$$\omega_{j,Mac,2017} = \sum_{k=1}^{k=11} \beta_k \omega_{jk,2012} + \epsilon_{j,2017}$$

  I then use only as my dependent variable PE_j, the residual of this prediction, understood as the Macron vote share "purged" from its predictable component based on the ideological composition of the local electorate in j:

$$PE_j = \epsilon_{j,\hat{2}017} = \omega_{j,Mac,2017} - \sum_{k=1}^{k=11} \hat{\beta}_k \omega_{jk,2012}$$

  The vote share of all candidates in the municipality in 2012 has significant predictive power, with the regression having an R^2 of 39 percent for the actual Macron vote share, but still preserves substantial variation in realized Macron support. The residualization implicitly neutralizes the impact of the Macron platform on the average municipality, thus only leaving room for heterogeneous effects depending on the initial tax burden.

3.3.4 Instrumental variables strategy

Even conditional on 2012 election results, and a wide array of observable characteristics X_j, omitted variables in $\bar{\epsilon}_j^C$, the city-level component of unobserved determinants of the Macron vote, may still drive any correlation between the pre-election housing tax exposure TH_j, on the one hand, and the residual Macron excess vote share adjusted for national 2012-2017 shifts, PS_j or PE_j. If populations with different political ideologies sorted themselves into towns with different housing values over the 2012-2017 period, a higher value of real estate in a municipality may still be associated with both the residual component of support for Emmanuel Macron, and the initial TH exposure.[20] Alternatively, local voters with idiosyn-

[20]It should be noted, however, that other components of the Macron platform, like restricting the French wealth tax (*ISF*) to the taxation of real estate, might have been expected to *reduce* the Macron vote share in high real estate price municipalities; and that adjusting for 2012 electoral behavior should offset most of this

cratic preference shocks for Emmanuel Macron over the 2012-2017 period may also have been more prone to vote for mayors and local officials who tend to set high local tax rates for the TH.

However, owing to the specific structure of the housing tax detailed in section 3.2.2, the tax base is constructed from 1970 appraisals that are, conditional on current home values, orthogonal to household sorting in 2017. Historical assessed values, VLC_j form an instrument for expected local benefits from the tax repeal, since they form a relevant but independent predictor of TH_j, that is uncorrelated with the non-tax local determinants of Macron vote shares, conditional on controls and 2012 electoral results. The two-stage least squares approach first predicts the average TH_j burden at the city level based on mean assessed values $V\bar{L}C_j$, current house values and current TH tax rates; it then regresses the Macron vote share on the predicted value of the tax $T\hat{H}_j^{IV}$, and the same set of controls. The first stage is:

$$T\hat{H}_j^{IV} = \alpha^{IV} + \delta\overline{VLC}_j + \lambda X_j + \nu_j \tag{3.3.4.1}$$

while the second-stage specification is the following:

$$\omega_{j,2017} = \tilde{\alpha} + \tilde{\beta}T\hat{H}_j^{IV} + \gamma X_j + \tilde{\eta}_j \tag{3.3.4.2}$$

As explained in section 3.2.2, the formulaic rental values established in the 1970s are correlated with current home values, since prices exhibit persistence over time. Thus, in the main IV specification, I control for current tax rates, current housing values, and their interaction, as well as a vector of controls including average receipts from other local taxes also based on the *VLC70*. This strategy also bypasses the other potential confounder, the correlation between local ideology and tax rates set by mayoral councils. Moreover, I use as my main outcome variable the "excess Macron vote share" (either PE_j or PS_j), constructed in the previous subsection, thus adjusting raw results for nationwide 2012-2017 shifts interacted with the composition of the local electorate. Even if 1970 VLC were associated with historical household sorting not taken into account by current home values, they are unlikely to be correlated with drivers of residual 2012-2017 changes in local voting decisions, other than

potential bias.

the housing tax burden.

3.4 Empirical results

3.4.1 Spatial variation in expected benefits and the Macron vote share

Baseline specification

To evidence the relevance and quantitative magnitude of the results in raw electoral data, I start from the simplest possible specification. I regress the vote share of Emmanuel Macron $\omega_{j,2017}$, in the first round of the French presidential election, on the average initial housing tax (TH) burden, TH_j, in municipality j. As described in section 3.2.3, I use data from all municipalities with available vote results in mainland France and Corsica.[21] In the baseline approach, I regress the Macron vote share in the first round of the 2017 election on average housing tax receipts per household in 2016 (corresponding to the last available year of housing tax assessment known to households before the April 23rd, 2017 election), as well as a number of controls at the municipality level. The main equation of interest is a cross-sectional regression across all 35,197 available municipalities in the sample, equation 3.3.2.2, where the preferred computation for the average burden in municipality j, TH_j (total tax receipts divided by the number of household with a tax return) is the dependent variable, and X_j is a (potentially empty) vector of controls at the municipality level. Throughout the paper, following Casaburi and Troiano (2016), I cluster standard errors at the province (*departement*) level, given the potential for spatial correlation in the error term. In subsection 3.4.3, I discuss the potential for spatial correlation and check the robustness of the results to using province fixed-effects and a spatial first differences design.

The results are presented graphically in figure 3.A.5a, while the estimated coefficients are reported in table 3.B.2. The baseline bivariate estimates in column 1 (displayed in figure 3.A.5a) show a strong, positive, economically large, and statistically significant correlation between the average housing tax burden in a municipality and the share of votes obtained

[21] French citizens living abroad are exempt from the TH; France's overseas provinces follow different rules in the setting of the housing tax and were already mostly exempt from it as of 2017; together, these groups excluded from my sample represent c. 6.6 percent of the overall electorate.

by Emmanuel Macron in the first round of the 2017 presidential election. The magnitude is substantial: a one standard deviation increase in the initial TH burden (c. EUR 250 or USD 280) is associated with a 1.7 percentage point higher Macron vote share. Going from the twenty-fifth percentile of initial housing tax burden to the seventy-fifth[22] is associated with a 2.46 percentage points higher Macron vote share in the first round of the presidential election.

As mentioned in section 3.3, the vote share for Emmanuel Macron is likely to be correlated with other features of the municipality's electorate, which, given the spatial sorting of households across space, could also be related to the initial housing tax burden. While the average TH was correlated with a number of characteristics of the municipality, there remained substantial heterogeneity in the housing tax burden even within narrowly defined strata. To adjust for the characteristics of the municipality, I include a series of municipality-level controls which might be correlated with the vote share for candidate Macron. Columns 2 and 3 include these additional controls. Column 2 includes only total population and mean income per capita (from the census and local public finance data sources). Column 3 includes, for a subset of larger municipalities and county seats, additional detailed controls from the *FiLoSoFi* database, on the structure of the local income distribution, notably the threshold of all income deciles, as well as the share of average local disposable income coming from capital income, pension income, and welfare benefits. Because these additional controls are only available for larger municipalities, they reduce the sample and shift its composition towards larger, richer towns, which tend to exhibit a lower response to the promised tax cut.[23] Nonetheless, all specifications tell a consistent story, and quantitative estimates of the tax burden effect remain large. The magnitude of the coefficient estimated by OLS is lower when including additional controls, with a EUR 1,000 increase in the average housing tax burden yielding a c. 4.45 percentage points increase in the Macron vote share, but the estimated effect remains highly significant and positive. Given the substantial variation in *commune* size, columns 4 to 6 repeat the specifications in columns 1 to 3, but re-weight observations by the number of registered voters in each municipality. Estimated coefficients

[22]Throughout the article, quantiles are re-weighted by the number of registered voters in each municipality.

[23]Section 3.4.2 explores evidence on this heterogeneous response in depth and provides some potential explanations.

are smaller than those for the corresponding un-weighted specifications, a reflection of the smaller estimated impact of the housing tax repeal in larger municipalities, where the TH represents a smaller share of disposable income and measurement error is likely to be larger due to variation across neighbourhoods in tax receipts.

Adjusting for local ideology

To disentangle the specific role of the tax cut proposal from unobserved, persistent determinants of local Macron support, I then adjust my electoral outcome measure (the Macron vote share in the first round) for two alternative proxies of the predicted Macron vote share, using the *polling shift* approach and the *prediction error* approach explained in section 3.3. This "ideology-adjusted" specification uses either PE_j or PS_j as the outcome variable, and regresses it, using ordinary least squares, on the preferred measure of the average housing tax burden.

$$\text{PS/PE}_{j,2017} = \alpha + \beta TH_j^1 + \gamma X_j + \eta_j \tag{3.4.1.1}$$

The results are displayed in figure 3.A.5b. The figure shows a tight, positive, and highly consistent relationship between the initial average housing tax per household and the deviation of the Macron vote share from its expected level based on the 2012 results of the five major presidential candidates. Table 3.B.3 summarizes the quantitative results of these ideology-adjusted specifications. Column 1 displays the results for the PE_j outcome, which amounts to controlling for the combined vote shares of each of the 12 candidates in the 2012 election. Columns 2 and 3 include additional controls, similar to the preceding subsection: column 2 includes widely available municipality level controls (population and mean income), while Column 3 includes additional detailed controls on the structure of the local income distribution, available only for the sub-sample of larger municipalities. Columns 4 to 6 repeat the same specifications, but for PS_j, the "polling shift" approach, as the dependent variable. The results demonstrate that, even conditional on the expected ideological composition of the local electorate, the Macron "excess vote share" in the first round of the 2017 election – whether one uses PE_j or PS_j as the outcome – is strongly and positively correlated with the average initial housing tax burden in a municipality. Coefficients in column

1 to 5 are all statistically significant at the 1 percent level, [24] and their magnitude is economically large. The results show that the "Macron excess vote share" increases by c. 0.5 percentage points in towns where the initial TH burden is one standard deviation higher, for my preferred estimate using the full sample, with controls, and PS_j as the dependent variable (column 5).

Instrumental variables specification

In this subsection, I turn to the two-stages least squares specification building upon the exclusion restriction assumption underlying the strategy outlined in section 3.3.4. *Conditional* on other local tax receipts for taxes based on the VLC, and on current home values, current TH tax rates, and their interaction, the average rental values (VLC) in a municipality matter for the definition of the TH tax base, but should not affect the local Macron vote share through any other channel than the expected benefits from the housing tax repeal. Thus, in the context of the election, any observed correlation between the VLCs and the Macron excess vote share can only be driven by the electoral impact of the campaign promise, validating the self-interested voter hypothesis.

Reduced-form results from this instrumental variable strategy are displayed in figure 3.A.6. It evidences the relationship (conditional on partialling out controls including current home values, TH tax rates, and their interaction) between the excess Macron vote share – using the polling shift adjustment method – and the average VLC in the municipality. The reduced-form results display a quantitatively substantial and tight relationship between the average rental value (derived from the 1970 assessments) in a town, and local support for Emmanuel Macron in excess of its predictable level based on the 2012 presidential election results.

Detailed quantitative results from this instrumental variable strategy are reported in table 3.B.4. Columns 1 to 3 use the raw Macron vote share $Macron_j$, as the outcome variable, instrumented by the mean VLC in the town. Column 1 only includes as controls other local tax receipts for taxes based on the VLC (the property tax, or *taxe fonciere sur les proprietes baties*, and the land-value tax, or *taxe fonciere sur les proprietes non baties*), as well as current median home values in the municipality, current TH tax rates, and their interaction. Column 2 in-

[24]The p-value for the coefficient in column 6 is 0.020, significant at the 5 percent level, with standard errors conservatively clustered at the province level.

cludes the additional controls (mean and median income and population) available for – almost – the entire sample, while Column 3 includes the set of additional controls available for larger municipalities and county seats already described in the previous section. Columns 4 to 6 repeat the same specifications, but using PS_j, the excess vote share adjusted for baseline pre-announcement nationwide polling shifts, as the outcome variable.[25] Across all specifications, I find strongly positive and statistically significant effects of the instrumented housing tax burden on the Macron vote share. According to the preferred estimates in column 5 (using the polling shift approach in the full sample with controls), a one standard deviation increase in the initial TH burden (instrumented by the 1970-determined assessed rental values, and conditional on current home values and tax rates) leads to a c. one percentage point higher excess Macron vote share. The magnitudes of the point estimates are larger using the instrumental variable strategy than the corresponding effects estimated via ordinary least squares, suggesting that omitted variable bias might attenuate the electoral effect in the baseline OLS regressions. One possible explanation is that locations with a higher concentration of profitable companies derive more revenue from the municipal corporate tax, and can therefore afford to vote for lower individual housing tax rates (and therefore average tax *bills*) for a given average housing tax *base*. If these economically dynamic cities also tend to host a higher share of centrist Macron voters, the OLS-estimated correlation between the initial housing tax burden and the Macron vote would be biased downwards relative to the actual causal impact of the tax repeal.

Vote-stealing effects To further investigate Macron's advantage in cities with a higher TH burden, I then turn to the differential impact of the initial TH on various other candidates in the 2017 election. The additional Macron support comes at the expense either of votes for other candidates, or through variation in turnout. Table 3.B.5 repeats the instrumental variables specification of table 3.B.4, using as the dependent variable the prediction error outcome for each of the other three major candidates (M. Le Pen, F. Fillon, J-L. Melenchon), as well as the abstention share (one minus turnout). The specification uses the prediction error approach to adjust for ideological composition, instruments for the TH burden using

[25] The number of observations is reduced relative to the OLS strategy in tables 3.B.2 and 3.B.3, since home values are derived from the DVF database, which does not include three Eastern France provinces representing close to 1,800 towns – see appendix 3.C.

the mean VLC in the municipality, and includes all available controls in the full sample. This exercise suggests that most of the impact of the TH on vote shares comes at the expense of support for far-right candidate Marine Le Pen, and, to a lesser extent, that a higher initial TH burden led to a small increase in abstention, with no statistically significant impact on support for other candidates.

3.4.2 Evidence on heterogeneous effects

The role of home-ownership The main effect of repealing of the housing tax, a reform fully financed initially by higher inter-governmental transfers from the State to local towns, was the implicit lump-sum expected redistribution from low- to high-TH burden municipalities. However, if local taxes are partly or fully capitalized into housing prices, an additional induced effect of the housing tax cut would be a windfall capital gain for homeowners. Such capitalization effects would imply a stronger electoral response of the Macron vote share to the initial housing tax burden in locations with a higher share of homeowners.

Hence, I investigate next the heterogeneous role of the initial housing tax burden on the vote between municipalities with varying shares of home-owners versus tenants. Unlike the property tax (*taxe fonciere*), the burden of the *taxe d'habitation* in France falls upon both home-owners and renters. Given the long-run nature of rental contracts and the well-established stickiness in nominal rents, economic incidence may closely follow statutory incidence.[26] Thus both types of households, if they expect to stay in their current town for some period of time, should be expected to react to the tax when making electoral decisions. Nonetheless, any capitalization effects of lower housing taxes into property prices would exclusively benefit homeowners, whose assets would increase in value, generating a windfall capital gain when selling their house. Therefore, the overall net benefit of the Macron proposed reform is expected to be higher for households who own their home than for those who rent it, consistent with the so-called "homevoter hypothesis" (Fischel, 2009).

In the spirit of Ahlfeldt and Maennig (2015), I re-run the ideology-adjusted specifications (equation 3.4.1.1), adding an interaction term between the share of homeowners in the municipality and the initial TH burden. I use data provided by INSEE and drawing from the full

[26]Standard rent contracts in France last three years. Moreover, rents cannot be adjusted upwards by more than a nationwide fixed indexation coefficient every year in the case of continuing rental relationships.

population census (*Recensement de la population*), providing me with the share of homeowners among primary residences in each municipality as of 2016. The baseline specification is:

$$\text{MacronShare}_{j,2017} = \alpha + \beta_T TH^1_j + \beta_O \text{Home-ownership}_{j,2016} + \delta \text{Home-ownership}_{j,2016} \times TH^1_j + \gamma X_j + \eta_j \tag{3.4.2.1}$$

where the coefficient of interest on the interaction term is δ.

Table 3.B.6 summarizes the results. I find that the interaction has a strong, positive, and statistically significant effect on the vote share of candidate Emmanuel Macron. Column 1 uses the baseline OLS specification with only the Macron vote share as a dependent variable and no additional controls beyond home-ownership. Columns 2 to 6 replicate the robustness tests already implemented for the baseline regression: instrumenting both the TH burden and the interaction term by the average VLC in the municipality (and its interaction with homeownership) while controlling for pre-period home values interacted with local tax rates (column 2), introducing municipality-level controls X_j, in the full sample (column 3), and repeating the previous three specifications using the Macron excess vote share as a dependent variable (using the polling shift approach PS_j) in columns 4 to 6. In all specifications, the additional marginal effect of home-ownership interacted with the TH burden remains economically large, and in all cases but column 6, significant at the five percent level.

To gauge the magnitude of the estimates of such heterogeneous effects, preferred results in columns 5 suggest that for cities at the first decile of home-ownership rates, a one standard-deviation increase in the housing tax burden in 2016 had an 1.27 percentage point effect on the excess vote share of candidate Macron; but that for those at the ninetieth percentile of home-ownership rates, the marginal effect of a one standard deviation increase on the excess vote share was 1.68 percentage points.

The role of exemptions As described in section 3.2, some households were exempt from the housing tax even before the Macron reform (depending on characteristics including household income, age, and the number of dependents). Therefore, we expect the impact of the proposed repeal on the Macron vote share to be smaller in towns where the share of exempt households is larger and where the housing tax is therefore a less salient issue and its cancellation a less valuable benefit.

To test for this heterogeneous effect, I re-run the ideology-adjusted specifications (equation 3.4.1.1), adding an interaction term between the share of exempt households in the municipality and the initial TH burden. The baseline specification is:

$$\text{MacronShare}_{j,2017} = \alpha + \beta_T TH_j^1 + \beta_O \text{Exempt share}_{j,2016} + \delta \text{Exempt share}_{j,2016} \times TH_j^1 + \gamma X_j + \eta_j \tag{3.4.2.2}$$

Table 3.B.7 summarizes the results. I find that the interaction between the initial TH burden and the share of households exempt from the tax has a strong, negative, and statistically significant effect on the vote share of candidate Emmanuel Macron, controlling for the direct effect of both variables. Column 1 uses the baseline OLS specification with the raw Macron vote share as a dependent variable and no additional controls beyond the share of exempt households in the municipality. Columns 2 to 4 replicate the robustness tests already implemented for the home-ownership interaction: introducing municipality-level mean income and population as controls X_j (in column 2), and instrumenting both the TH burden and the interaction term by the mean assessed value in the municipality (and its interaction with the exempted share), while using the Macron excess vote share as a dependent variable (using the polling shift approach PS_j), in columns 3 and 4. In all specifications, the additional marginal effect of the share of exempted households interacted with the TH burden remains negative and economically large.

The role of town size As mentioned above in section 3.4.1, the effect size in ordinary least squares regressions appears smaller when restricting the sample to larger municipalities. Column 1 of Table 3.B.8 shows this heterogeneous effect size more directly, using the raw Macron vote share as an outcome, and suggests that the interaction of population size and initial average TH burden has a negative coefficient estimate. More populated towns appear less responsive to the TH repeal promise. Column 2 shows that this heterogeneous effect also appears when using the PS_j polling shift dependent variable as an outcome.

There are, however, three plausible explanations for the lower estimated effect using OLS in larger cities. First, a potential reason for the lower estimated effects in larger municipalities may simply be measurement error. Indeed, averaging at the municipality level is likely to lead to more attenuation bias in a highly populated and diverse city of 500,000 inhabitants than in a town with a population of 500. Because larger towns have a wider diversity of neighbourhoods and quality categories, averaging the TH burden at the municipality level creates measurement error in the independent source of variation, biasing the effect towards zero in this subset of towns.

Second, because larger towns are characterized by higher incomes on average, the share of the housing tax in the median household budget is smaller in these areas, making a given average euro amount less salient for residents of these cities. Column 3 of table 3.B.8 suggests that, when using as an independent variable the ratio of the income tax to median income in the municipality, one cannot reject the absence of heterogeneous effect size for larger municipalities; the same occurs when including directly mean and median income as controls in column 4.

Third, because larger towns often benefit from the presence of larger corporations and associated revenue from local business and corporate property tax, they are able to set lower rates on the politically more salient local residential housing tax, making the housing tax a less politically salient issue in these cities and a less relevant burden on a typical household budget. Column 5 of table 3.B.8 indeed shows that, when instrumenting the housing tax by local mean VLC, and the interaction term by the interaction of the mean VLC and population, the effect size for larger municipalities is again not statistically different than for smaller towns.

3.4.3 Robustness

Adjusting for spatial correlation The spatial correlation in both Macron excess vote shares and initial average housing tax, visible in figures 3.A.3a and 3.A.3b, suggests a need to account for regional or local shifts in electoral outcomes, potentially driven by unobserved correlates of the housing tax burden. Although all regressions already cluster standard errors at the province level to account for such spatial correlation, I follow two distinct empirical approaches to control more flexibly for regional shifts and clarify the distinct role of the initial TH burden.

The first strategy includes province fixed-effects in the baseline regressions, to account for any additional regional shifts in the support for Emmanuel Macron not accounted for by national polling shifts. This also removes systematic local variation in the TH burden, such as the updating of rental values in 1980 at the province level, or province-level funding needs shocks for local authorities which may have had spillovers on the tax rates set locally by municipalities. Table 3.B.9 repeats the specifications of table 3.B.4, but including such province fixed effects. The magnitude of the coefficients (using either the IV specification or the OLS, and either the raw vote share or the excess vote share adjusted for polling shifts) is slightly increased relative to the baseline measures, and all estimates remain significant at the 0.1 percent level.

My second strategy to control for local shocks uses a strategy akin to the spatial first-difference design (SFD) developed by Druckenmiller and Hsiang (2018). I use geographic adjacency files provided by OpenStreetMap to compare each municipality j to a weighted average of the set of immediately neighbouring municipalities $\mathcal{N}(j)$,[27] and compute for each variable W_j (the outcome variable, the initial tax exposure, the instrument, and the controls), a "spatial first difference" equivalent, $\Delta W_j = Y_j - \bar{W}_{N(j)} = Y_j - \sum_{k \in \mathcal{N}(j)} \frac{Pop_k}{\sum_{k' \in \mathcal{N}(j)} Pop_{k'}} Y_k$. I then regress the *difference* in electoral outcomes $\Delta Macron_j = Macron_j - Macron_{N(j)}$ (or $\Delta PS_j = PS_j - PS_{N(j)}$) between a city and its set of neighbors on the *difference* in initial TH burdens $\Delta TH^1_j = TH^1_j - TH^1_{N(j)}$, using either ordinary least squares or instrumental variables.

This spatial differencing strategy is conceptually close to a first-differences panel data regression over time, when dependent and independent variables are first subtracted their lagged values. As discussed in Druckenmiller and Hsiang (2018), the goal is to remove any remaining unobserved spatially correlated omitted variable bias. The baseline (ideology-adjusted) ordinary least squares specification, in equation 3.4.3.1, thus relates the difference in the Macron (excess) vote between two towns that share a geographic border, and their initial relative TH burden.

$$\Delta\text{MacronShare/PE/PS}_{j,2017} = \alpha + \beta \Delta TH^1_j + \gamma \Delta X_j + \eta_j \tag{3.4.3.1}$$

[27] I use a spatial first-difference design with respect to a population-weighted average of all neighbouring municipalities. Results are virtually unchanged when using only the largest neighbouring municipality.

I also run the corresponding instrumental variables regression, using as an instrument for ΔTH_j^1 the spatial first difference (SFD) in mean VLCs, ΔVLC_j, and controlling for the SFD of median home values and its interaction with the SFD of local tax rates, as well as the SFD of other local taxes based on the VLC. Table 3.B.10 describes the results. The magnitude of the coefficients is stable and significant at the 0.1 percent level across specifications, suggesting the removal of spatially correlated endogeneity indeed makes estimates more precise. IV coefficients are, as earlier, larger than OLS, but overall similar in magnitude to the baseline IV specifications.

Varying the definition of the average tax burden I compute alternative definitions for the municipality-level value of the housing tax burden in 2016, and show the robustness of the results to the use of these alternative measures in table 3.B.11. Using data from the *REI* file, I experiment with a series of alternative measures, including the ratio of total TH receipts to inhabitants (column 1), to registered voters (column 2); the ratio of the average tax payment to the mean income per consumption unit (column 3) and to median household income (column 4) in the municipality; and the euro-denominated value provided by the French Finance ministry publicly available document mentioned in section 3.2.3, published in 2018 after the Macron election, quantifying the benefit of the housing tax cut at the municipality level for the "average" household (column 5).

For ease of interpretation, independent variables are normalized to have a mean of zero and a standard deviation of one, so that the figure in each column denotes the impact of a one standard deviation increase in the TH burden measure on the Macron excess vote share, using the polling shift approach (PS_j), and instrumenting the variation by the mean 1970-determined VLC in the municipality. While coefficients vary in magnitude, given the varying definition of the housing tax burden, all imply a quantitatively large effect, consistent with the baseline results.

Placebo test: the 2012 election In spite of the instrumental variable strategy, and several robustness checks, it is not entirely inconceivable that omitted variable bias (or alternative channels correlated with assessed rental values, even conditional on current house prices and tax rates) could still lead to a spurious correlation between local TH payments (or local

VLC) and political alignment. To gauge the risk and potential magnitude of such a bias, I use the 2012 presidential election as a placebo test.

The *Taxe d'Habitation* was not a salient wedge issue during the 2012 election. None of the major candidates campaigned on its repeal, it was not mentioned in any of the televised debates, and only a few proposals to adapt the tax were made, mostly through income-based modulations (in the case of then Socialist candidate and future president Francois Hollande) or proposals to integrate it in the national progressive income tax (in the case of far-right candidate Marine Le Pen). All such proposed measures ranked relatively low on candidates' platforms, and did not generate any national news coverage, beyond specialized trade publications destined for local officeholders.

As in section 3.4.1, I purge the vote share of each of the main 2012 contenders from its predictable component by removing the systematic association with previous election results. This time, I use 2007 election results at the municipality-level to construct the $PE_{j,2012,k}$ "prediction error approach" variable, for k corresponding to each of the top five candidates in 2012, based on the same methodology as above. I regress the 2012 vote share of candidate k on the full set of vote shares (as a share of registered voters) of each of the 12 candidates (plus abstention) in the first round of the 2007 election.

$$\omega_{j,k,2013} = \sum_{l=1}^{l=13} \beta_l \omega_{jl,2007} + \epsilon_{j,2012}$$

I then use only the residual of this prediction, as my dependent variable $PE_{j,2012,k}$ for candidate k's vote share "purged" from predictable realized shifts due to the ideological composition of the local electorate in j:

$$PE_{j,2012,k} = \omega_{j,k,2012} - \sum_{l=1}^{l=13} \hat{\beta}_l \omega_{jl,2007}$$

This residualization amounts to adjusting flexibly in the main specification for the share of each candidate in the municipality in 2007.

I then regress the value of $PE_{j,2012,k}$, for each of the top five candidates, on the 2012 housing tax burden gathered from the *REI* 2012 database, as well as municipality-level controls provided by the FiLoSoFi 2012 data provided by INSEE. Table 3.B.12 repeats the ideology-adjusted specification (with controls for population and median income) for the top five candidates. As in all previous specifications, standard errors are clustered at the province level. Only one out of the top five contenders (Francois Bayrou) in 2012 exhibits a correlation of his excess vote share with the initial TH burden significant at the five percent level. The magnitude of that correlation in 2012 is at least one order of magnitude smaller than the effect size measured on the Macron excess vote share in 2017, and the R-squared of 0.003, despite the addition of controls, is negligible, and two orders of magnitude lower than the R-squared of the corresponding ideology-adjusted specification for the Macron excess vote share regressions in column 2 of table 3.B.3. Moreover, given that Bayrou endorsed Macron in the 2017 election, even assuming the potential existence of a quantitatively small omitted variable bias pushing voters *away* from centrist candidates in high TH-burden municipalities would lead to our main 2017 estimates being biased *towards zero*. At any rate, even if one cannot fully rule out the existence of omitted channels and variables biasing the central estimates, the absence of a statistically significant or quantitatively relevant impact of the TH burden in the previous election provides strong evidence that the TH itself is not conditionally correlated to persistent, unobservable drivers of ideological alignment at the local level, and that its 2017 effect is indeed a reflection of the self-interested voter hypothesis.

3.5 Aggregate impact

In this final section, I provide evidence, using aggregate Internet searches, polling, and prediction market data, that Macron's promise to repeal the housing tax indeed coincided with a rise in his appeal in the broader electorate. Moreover, I show that back-of-the-envelope calculations suggest the quantitative magnitudes of these aggregate effects are consistent with the range of implied effects from my spatial heterogeneity estimates.

3.5.1 Internet searches

Taxpayers use Internet search engines to look for information about their expected tax liability, and tend to do so around particular events, which can be interpreted as a form of rational attention, or a form of salience-based updating. Using Google Trends, Hoopes, Reck, and Slemrod (2015) have documented that taxpayers in the United States search for tax information online especially more around trigger events, and that presidential elections, in particular, are accompanied by increased search effort, notably when candidates make taxes a salient dimension of their policy platform.

Macron announced the policy on February 24th, 2017. Using monthly data from Google Trends's search engine,[28] I first show, in figure 3.A.7, that searches for the housing tax in 2017 and 2018 broke with their regular seasonal pattern (which usually exhibits a substantial spike in October-November, the deadline for payment), reflecting the increased salience of the housing tax during the 2017 presidential campaign. The abnormal attention drawn to the housing tax started with the Macron announcement in February 2017 and spiked in May 2017, the date of the runoff stage of the election. Searches related to the housing tax itself were particularly prominent in the months leading to the first round of the presidential election. They then rose again in July 2017, after a keynote address by newly nominated Prime Minister Edouard Philippe suggested the promised reform might be delayed, and a rift ensued between Macron and part of his cabinet.

Then, again using monthly data, I include as keywords joint searches for the centrist candidate and the housing tax ("taxe habitation Macron"). As shown in figure 3.A.8, such joint searches rose fast in the period immediately following the announcement, and remained elevated throughout the campaign, peaking in May, the month of the run-off stage, but reaching high levels even in the following months, notably around and after the July 2017 policy address and subsequent conflict. In figure 3.A.9, looking at more granular daily data, by restricting the sample to the year 2017, I show that searches mentioning both Macron and the TH rose fast on the day of the announcement, suggesting widespread interest in the reform. They also rose in the runup to the first round of the election, and in the days following the presidential debate between E. Macron and M. Le Pen in between the two rounds, when the housing tax was one of the salient issues discussed.[29]

[28] Hoopes, Reck, and Slemrod (2015) provide detailed information on the computation of Google's index for the "propensity to search", a 0 to 100 index where 100 corresponds to the highest relative value of searches for a term over the specified sample period in a given area. I restrict the sample to France over a ten year period, from January 1, 2010 to January 1, 2020. The results do not depend on the time window used.

[29] See e.g. Ouest-France, May 4th 2017, *Débat présidentiel. Ce qu'il faut retenir de l'affrontement entre Macron et Le Pen*

3.5.2 Prediction markets

Prediction markets aggregate individual assessments of a race in progress, and, due to the incentives of participants to make accurate forecasts, they can provide high frequency market-based estimates of a candidate's probability of winning an election. While such data suffer from known limitations (Wolfers and Zitzewitz, 2004), especially for low-liquidity markets, they can help identify the *ex ante* likely impact of events on electoral outcomes at relatively high-frequency.[30] I use data retrieved directly from PredictIt,[31] the most widely used prediction market, in a 30-day window around the reform announcement, on the daily closing price in a contract paying a dollar in case of a Macron win in the election. Relative to political prediction markets for the United States, French markets have much lower volume, given their lesser penetration and social acceptability in the wider French society (Charpentier, 2017). Results displayed in figure 3.A.11, graphically show that, in a thirty day window around the announcements, Macron's predicted winning probability, as estimating from the closing price of a contract paying a dollar in the case of a Macron victory substantially in the days immediately following the announcement. The fifteen-day means immediately before and after the announcement exhibit a more than ten percentage points difference in estimated winning probabilities. One should note, however, that, two days before the policy announcement, on February 22$^{\text{nd}}$, 2017, centrist politician and 2012 presidential candidate Francois Bayrou announced his support for E. Macron. However, the sustained rise in Macron's chances immediately after the TH repeal announcement, over and beyond the initial gain from the announcement of Bayrou's support two days earlier, is likely to reflect market estimates of the effect of the promised housing tax cut, especially so since Bayrou's support would already have been at least partially priced in Macron's victory chances.

3.5.3 Daily polling results

To provide evidence that the promise to exempt a substantial share of the electorate from the housing tax had a significant and immediate effect on voting intentions for Emmanuel

[30] See Coulomb and Sangnier (2014) for an application of the use of such data in the French case in the 2007 presidential election, ten years earlier.

[31] PredictIt data are directly available for download from the following page: https://www.predictit.org/markets/detail/2947/Who-will-win-the-most-votes-in-the-first-round-French-presidential-election-in-2017.

Macron, I use data from daily polls released before and after the announcement. Some pollsters have assembled anecdotal evidence around Election day that the housing tax cut was indeed an emblematic and salient proposal of the Macron campaign. [32] In particular, I use data provided by PollsPosition, a French polls aggregator.[33] These data aggregate all national polls realized in France using a representative sample and a method agreed upon by the National Polling Commission and defined in a July 1977 law. They include data from eight different pollsters. Polls are dated by the median fieldwork date. Using an event study design, I measure predicted voting intentions for Emmanuel Macron around the announcement of the policy. As shown in figure 3.A.10, Macron's support (as measured by voting intentions in the adult population for the first round of the 2017 presidential election) rose durably in the days immediately following his morning show announcement of his intention to scrap the housing tax. Akin to the observed prediction market impact around the announcement date, the rise could partly stem from the endorsement received by Macron from centrist 2012 candidate Francois Bayrou on February 22$^{\text{nd}}$. The limitation to daily data in the event study does not allow me to fully disentangle the direct impact of Francois Bayrou's endorsement from the impact of the proposed housing tax repeal, especially given that all polls used in the specification take between two and three days of fieldwork to be conducted, thus confounding the estimates of the relative impact of the TH tax cut and the Bayrou endorsement. However, the fact that the placebo analysis in section 3.4.3 shows no local correlation between the 2012 vote for Bayrou and the initial TH burden, and the coincidence with online searches as a proxy for the salience of the TH issue, both point towards the TH repeal pledge being the key driver of the rise in polling intentions.

Taking the estimates of the effect of the pledge obtained from the instrumental variables specification in section 3.4.1, one can perform a simple aggregation exercise to check their

[32] See notably the IFOP Report n. 172, in November 2017, *"L'exonération de la taxe d'habitation : mesure totémique du candidat-président Macron"*. According to a Harris Interactive April 20$^{\text{th}}$ 2017 poll, two days before the first round of the presidential election, among 1022 respondents, a plurality (55 percent) cited the TH repeal as the most memorable and convincing policy platform of the campaign; 72 percent of those intending to vote for Emmanuel Macron cited the policy as the most convincing of the campaign. Another July 2017 Harris Interactive poll, which sampled 978 individuals soon after the presidential election, the repeal of the TH was Macron's most favored proposal among those tested: 80 percent of respondents favored the policy.

[33] I am thankful to Berengere Patault and Alexandre Andorra for sharing the PollsPosition data. These data are now made publicly accessible at https://www.pollsposition.com/home.

consistency with high-frequency polling outcomes. Assuming a linear relationship between the Macron vote share (as a percentage of registered voters in a municipality) and the initial TH burden, the promise was enough to receive an additional number of votes equal to $\hat{\beta}$ times the weighted average housing tax across municipalities (where the weights N_j are the number of registered voters in the municipality). Using the main IV estimate of column 5 in table 3.B.4, $\hat{\beta} = 4.07$, implies that the aggregate number V of additional votes received by Emmanuel Macron a a result of the campaign promise was roughly equal to $V = \sum \hat{\beta} \times TH_j \times N_j$, which, given that the (weighted) mean TH burden was around EUR 760, corresponds to an additional 1.4 million votes (3.1 percent of registered voters, or 3.75 percent of votes cast). Such an electoral boost is consistent with the approximate 3.5 percentage points increase in average polls results around the announcement date, evidenced in figure 3.A.10; it would have been enough to push Macron to reach the runoff stage.

3.6 Concluding remarks

Numerous studies have documented the retrospective role of economic conditions or transfers in triggering support for incumbents. Researchers often rationalize these findings by voters updating imperfect prior beliefs about candidate's trustworthiness or competence. On the other hand, despite self-interested voting being a tenet of political economy theory, there is surprisingly little empirical evidence on the causal role of forward-looking promises for electoral outcomes. This paper contributes to an emerging literature causally studying the "prospective" voting motive, and in particular, the electoral impact of promised tax cuts. I exploit the housing tax repeal proposed by Emmanuel Macron in his 2017 French presidential bid to provide local and aggregate evidence that the electorate responds to monetary campaign promises by candidates, enough to sway a major election in a large, advanced economy. Towns with higher expected tax savings, as predicted by historical assessed home values from the 1970s still in use as of 2017, and adjusting for both current housing values and tax rates, witnessed substantially larger support for Emmanuel Macron in the first round of the presidential election. Controlling for local demographic characteristics, income distribution patterns, or past electoral outcomes in the municipality, does not eliminate the estimated impact of the initial housing tax burden. Even when restricting my estimation strategy to variation between neighbouring municipalities, the results evidence an economically and statistically significant impact of the predicted local tax cut on support for the candidate pledging it. Consistent with the causal mechanism, heterogeneity analyses suggest an especially strong impact of the promised tax cut in municipalities with a larger share of homeowners, a smaller share of exempt households, and a lower share of secondary residences.

To pinpoint the causal impact of material promises on voting behavior, I avail myself of a nationwide policy platform differentially affecting the burden of local taxation. This feature generates substantial divergence across fine geographic areas in exposure to the tax cut – even among contiguous municipalities, and conditional on observables likely to affect the performance of each presidential candidate. Moreover, the nature of the French housing tax, owing to the pre-determined and formulaic nature of the tax base assessment, allows for a plausible instrumental variable strategy to deal with potential endogeneity concerns. As demonstrated by opinion polls and Internet searches, the policy was clearly identified by voters as a key plank of the centrist candidate's platform, especially given the high salience of the housing tax in the electorate. Because of the absence of an incumbent, voters had relatively little information on any of the contenders' potential skills as head of State. They were therefore unlikely to vote based on retrospective evaluation considerations, and more sensitive to prospective policy concerns embodied in candidates' platforms.

The efficacy of material campaign promises in delivering votes remains a contentious topic, both among social scientists and practitioners of politics. The self-interested voter hypothesis has recently lost some of its centrality in academic studies of voting behavior. Indeed, it has been suggested, that, as living standards improved in the long-run, narrow material interests lost their pre-eminence in driving political behavior. Thomas Frank's widely acclaimed 2007 book, *What's the matter with Kansas?* (Frank, 2007), forcefully argued that alternative determinants of the vote, from "cultural anxiety" to altruistic considerations, accounted for the apparent disconnect between voters' economic self-interest and their political preferences. This paper suggests that, even at times of high polarization along cultural or identity concerns, costly but credible and salient tax cuts pledges may be enough to win the upper hand in major electoral contests, and that expected material gains and losses from a policy platform are still relevant determinants of individual voting behavior.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendices

3.A Figures for Chapter 3

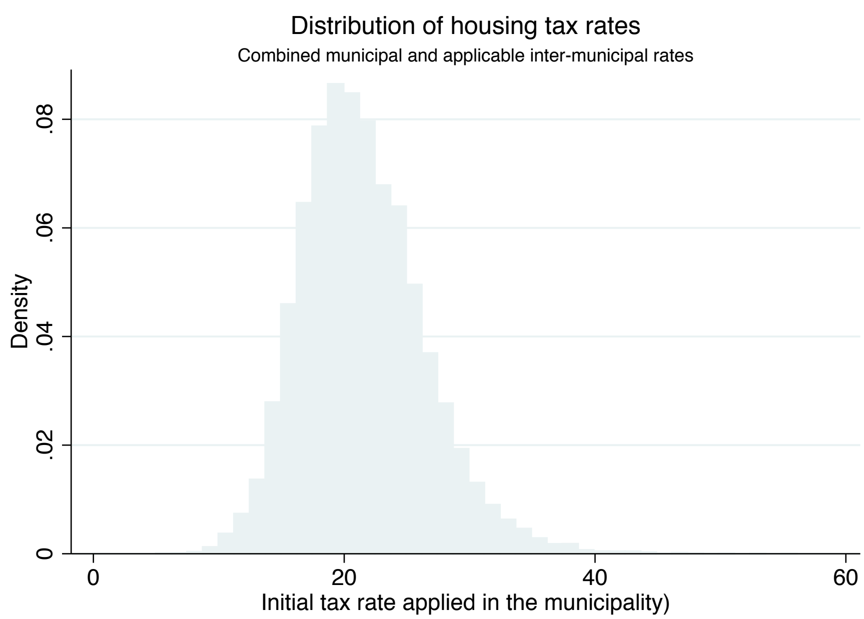

Figure 3.A.1: **Housing tax rate distribution**

The figure is a histogram plotting the frequency of initial housing tax rates by municipality, combining the rate of the municipality and the inter-municipal cooperation agreement.

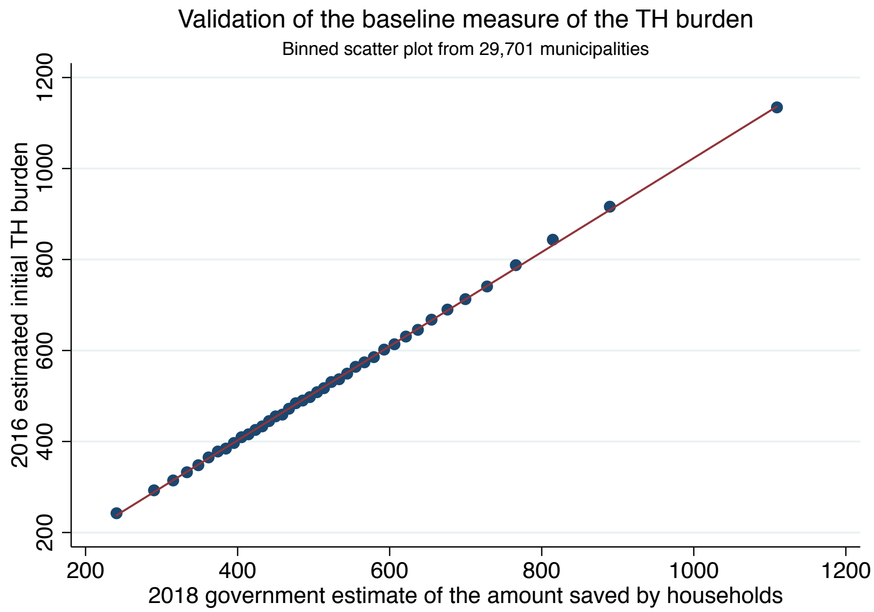

Figure 3.A.2: **Validation of the housing tax burden measure**

The figure is a binned scatter plot plotting the baseline estimate of the initial housing tax burden per household in a municipality, in euros, against a government estimate of average savings per household published in 2018, after the reform had been implemented. The x-axis of each dot is the mean value of the government estimate in the percentile. The y-axis is the average value of the tax burden (baseline measure) in the percentile. The line plots the predicted values from a linear regression model.

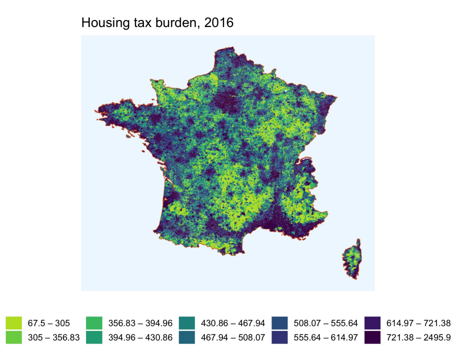

(a) Map of the housing tax burden

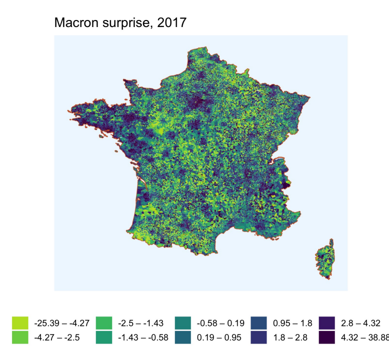

(b) Map of the Macron polling shift surprise

Figure 3.A.3: **The initial TH burden and the Macron excess vote share**

The left panel maps the initial housing tax burden per household in a municipality, in euros, as of 2016, the last year when the value of the tax was known to households before the presidential election. Darker values denote higher values of the initial TH burden. The detailed computation of the average housing tax burden is explained in section 3.2.3. The right panel maps the excess vote share is computed using the "polling shift" approach which adjusts for nationwide swings from 2012 candidates to Emmanuel Macron in polls immediately preceding the week of the announcement of the housing tax repeal. Darker values denote higher values of the excess vote share, the detailed computation of which is explained in section 3.4.1

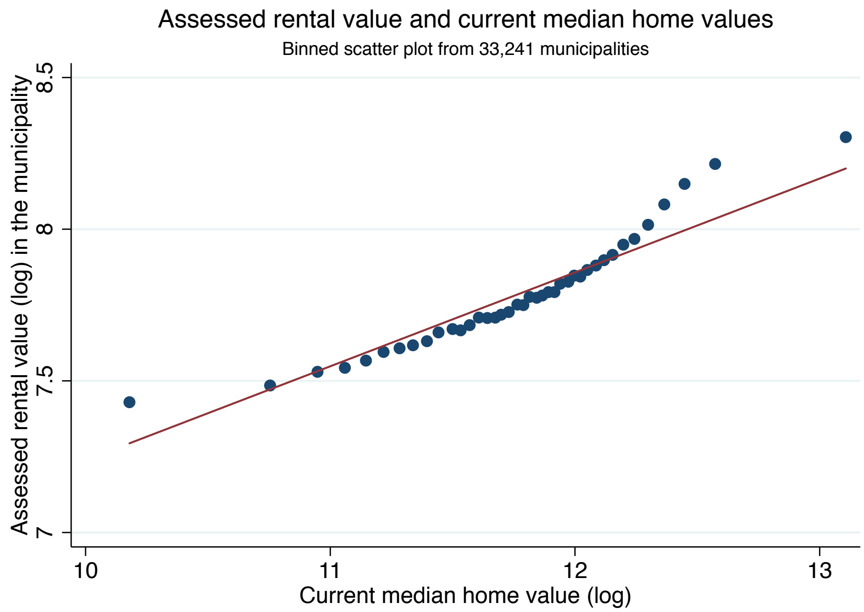

Figure 3.A.4: **Validation of the market value measure**

The figure is a binned scatter plot plotting the baseline estimate of the mean assessed rental value (VLC) in the municipality, against an estimated median home value in the municipality over the period 2014-2019 from the *Demande de Valeurs Foncieres* database. The x-axis of each dot is the mean value of the housing market value in the percentile. The y-axis is the average value of the municipality-level VLC in the percentile. The line plots the predicted values from a linear regression model.

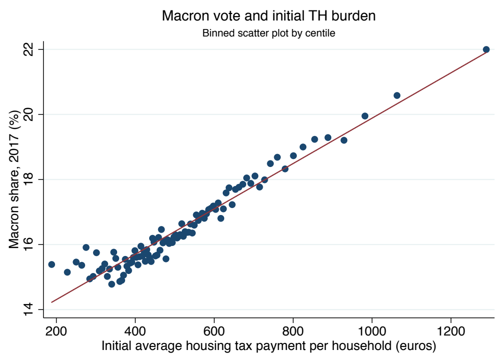

(a) Macron first-round vote share and housing tax

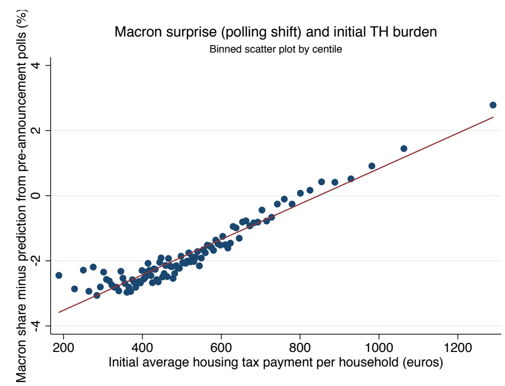

(b) Macron first-round vote surprise and housing tax

Figure 3.A.5: **Baseline and ideology-adjusted specification**

The left panel is a binned scatter plot plotting Macron's first round vote share in the 2017 French presidential election against the initial housing tax burden per household in a municipality, in euros. The right panel plots Macron's first round excess vote share against the initial housing tax burden. The excess vote share is computed using the "polling shift" approach which adjusts for nationwide swings from 2012 candidates to Emmanuel Macron in polls immediately preceding the week of the announcement of the housing tax repeal. The x-axis is partitioned into percentiles. The x-axis of each dot is the mean value of the housing tax burden in the percentile. The y-axis is the average value of the Macron (resp. excess) vote share in the percentile. The line plots the predicted values from a linear regression model.

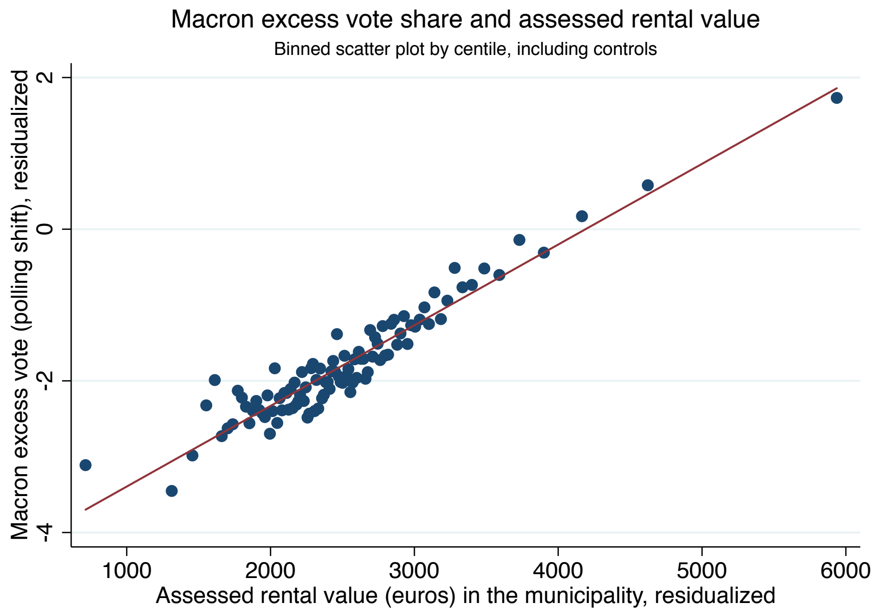

Figure 3.A.6: **Macron first-round vote surprise and initial VLC: reduced form**

The figure is a binned scatter plot plotting Macron's first round excess vote share against the initial average 1970-determined rental value in a municipality, in euros. The excess vote share is computed using the "polling shift" approach which adjusts for nationwide swings from 2012 candidates to Emmanuel Macron in polls immediately preceding the week of the announcement of the housing tax repeal. The x-axis is partitioned into percentiles. The x-axis of each dot is the mean value of the rental value in the percentile. The y-axis is the average value of the excess vote share in the percentile. The line plots the predicted values from a linear regression model.

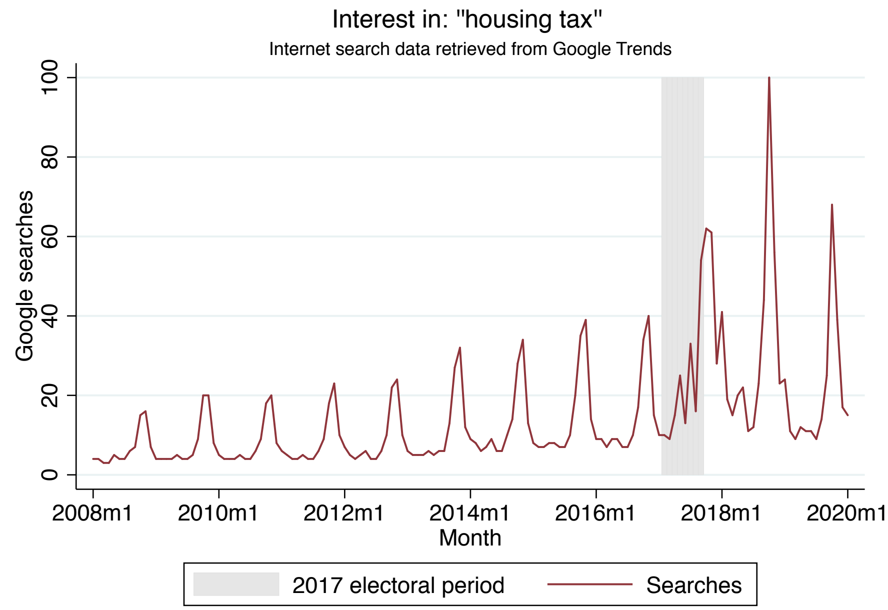

Figure 3.A.7: **Google searches for the housing tax - 2008-2020**

The figure plots monthly values from Google Trends for searches containing the term "taxe habitation" in France over the period 2008-2020. The index denotes Google's computation of a propensity to save, normalized to 100 at its maximum over the period. The shaded area denotes the period starting with the Macron repeal announcement, and ending with the finalization of the conflict between Macron and his cabinet over the implementation of the reform.

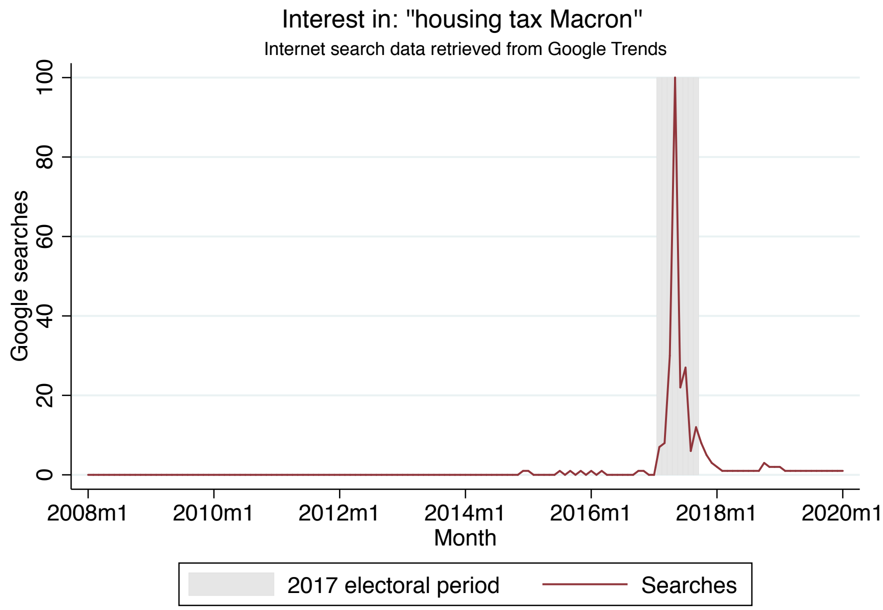

Figure 3.A.8: **Google joint searches for "housing tax" and "Macron"**

The figure plots monthly values from Google Trends for searches containing the terms "taxe habitation Macron" in France over the period 2008-2020. The index denotes Google's computation of a propensity to save, normalized to 100 at its maximum over the period. The shaded area denotes the period starting with the Macron repeal announcement, and ending with the finalization of the conflict between Macron and his cabinet over the implementation of the reform.

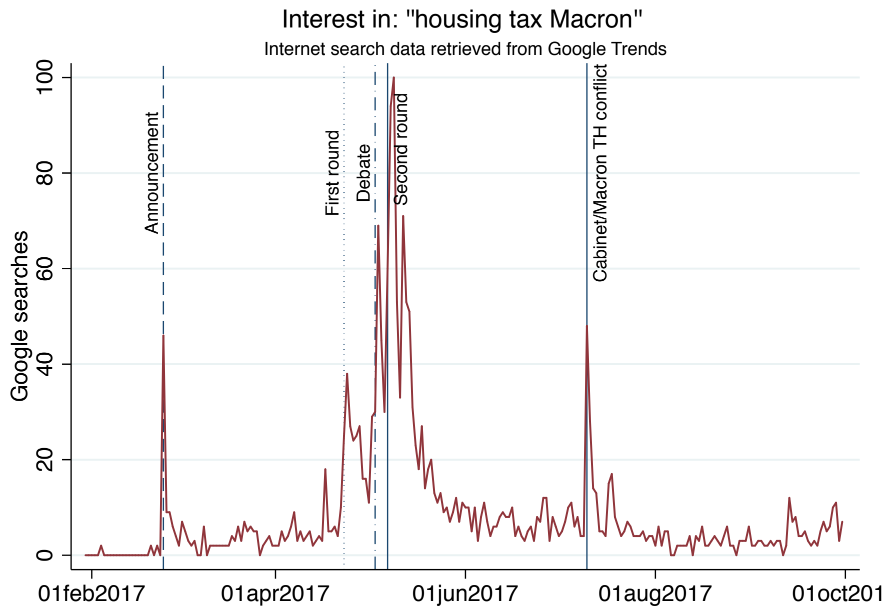

Figure 3.A.9: **Google joint searches for "housing tax" and "Macron" - campaign period**

The figure plots daily values from Google Trends for searches containing the terms "taxe habitation Macron" in France over the period January to September 2017. The index denotes Google's computation of a propensity to save, normalized to 100 at its maximum over the period. The lines denote successively the Macron repeal announcement, the first round of the presidential election, the presidential debate, the runoff stage, and the height of the conflict between Macron and his cabinet over the implementation of the reform.

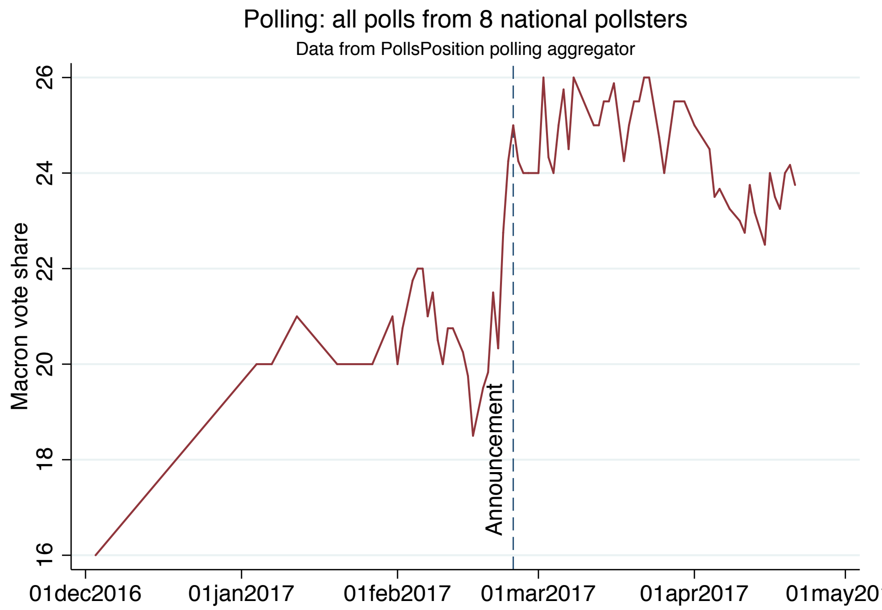

Figure 3.A.10: **Polls: first round voting intentions for E. Macron**

The figure plots voting intentions for E. Macron in the first round, using all polls aggregated by PollsPosition. The vertical line denotes the Macron repeal announcement.

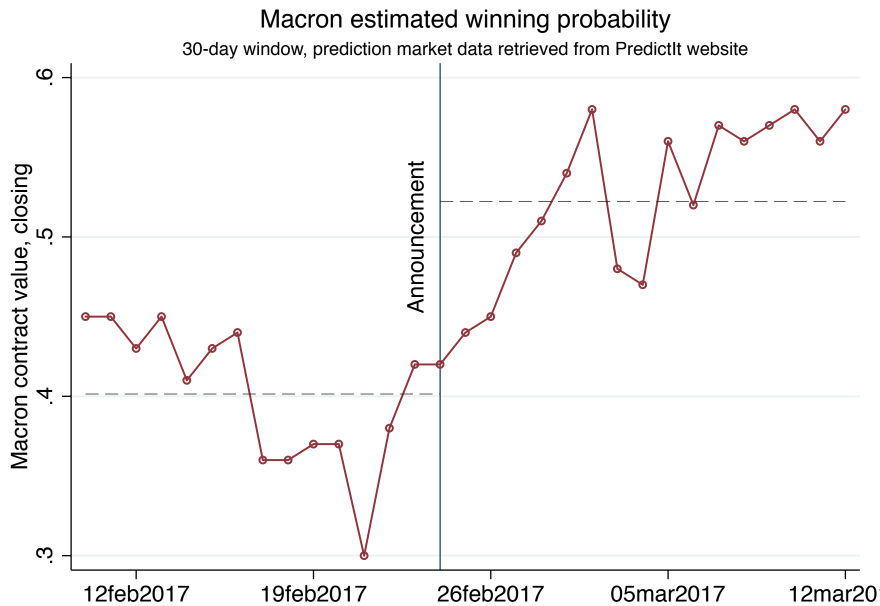

Figure 3.A.11: **PredictIt estimated winning probability of Macron**

The figure plots daily values from PredictIt for the closing price of a contract paying one dollar in the case of a Macron final victory in the French presidential election in a thirty day window around the announcement. The vertical line denotes the Macron repeal announcement; the horizontal line denote the pre- and post-mean for the contract value in the fifteen days periods immediately preceding and following the announcement.

3.B Tables for Chapter 3

Table 3.B.1: Summary statistics

	(1)				
	Mean	Std. dev	Min	Max	N
Population	1920.7	15244.94	0.0	2243739.0	35502
Registered voters	1259.6	8841.56	6.0	1301637.0	35277
Macron share of registered votes	16.5	5.02	0.0	66.0	35277
Initial TH burden	518.5	185.43	79.1	2609.0	35197
Mean VLC in the municipality	2527.8	860.51	0.0	12571.0	35212
Median Income per household	20892.6	2986.77	10932.0	48288.1	31789
Mean income per capita	13117.8	3275.43	0.0	69642.8	35502
Threshold of first decile of income	12274.2	1963.09	5758.8	20432.0	5279
Threshold of ninth decile of income	36453.2	8429.33	22528.3	128772.0	5279

Data come from a variety of sources, including the French Interior Ministry, French national statistical institute, French Finance Ministry, French Territorial Planning Authority. The number of available observations for some detailed income distribution variables is lower, reflecting statistical secrecy rules.

Table 3.B.2: Impact of the TH on Macron vote share: baseline specification

	(1)	(2)	(3)	(4)	(5)	(6)
Initial TH burden (EUR '000s)	6.97****	5.14****	4.45****	5.13****	2.12*	2.83*
	(0.85)	(0.99)	(1.18)	(1.34)	(1.18)	(1.60)
Town size weights	No	No	No	Yes	Yes	Yes
Full-sample controls	No	Yes	Yes	No	Yes	Yes
Restricted sample controls	No	No	Yes	No	No	Yes
R-Square	0.07	0.08	0.41	0.07	0.35	0.52
Observations	35197	35197	5207	35197	35197	5207
Clusters	96	96	96	96	96	96

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$.

Standard errors in parentheses, clustered at the province level. Columns 1 to 3 are unweighted ordinary least squares regressions; Columns 4 to 6 are weighted by the size of the local electorate. Columns 2 and 5 include as controls mean and median income, as well as population size. Columns 3 and 6, in addition to the same controls, include all deciles of income and the share of disposable income coming from pensions, capital income, and social transfers.

Table 3.B.3: Impact of the TH on Macron vote share: ideology-adjusted specification

	(1)	(2)	(3)	(4)	(5)	(6)
	"Prediction error" approach			"Polling shift" approach		
Initial TH burden (EUR '000s)	3.32****	1.18***	1.57**	4.34****	2.18****	1.93**
	(0.67)	(0.44)	(0.60)	(0.78)	(0.60)	(0.78)
Full-sample controls	No	Yes	Yes	No	Yes	Yes
Restricted sample controls	No	No	Yes	No	No	Yes
R-Square	0.05	0.08	0.48	0.11	0.45	0.62
Observations	35197	35197	5207	35197	35197	5207
Clusters	96	96	96	96	96	96

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 to 3 are ordinary least squares regressions for the PE_j or "prediction error" dependent variable; Columns 4 to 6 are ordinary least squares regressions for the PS_j or "polling shift" variable. Columns 2 and 5 include as controls mean and median income, as well as population size. Columns 3 and 6, in addition to the same controls, include all deciles of income and the share of disposable income coming from pensions, capital income, and social transfers.

Table 3.B.4: Impact of the TH on Macron vote share: instrumental variables specification

	(1)	(2)	(3)	(4)	(5)	(6)
	Raw Macron vote share			"Polling shift" excess vote share		
Initial TH burden (EUR ’000s)	7.74****	6.09****	4.85****	5.62****	4.07****	2.67****
	(0.99)	(1.20)	(1.25)	(0.53)	(0.64)	(0.62)
Local tax controls	Yes	Yes	Yes	Yes	Yes	Yes
House value controls	Yes	Yes	Yes	Yes	Yes	Yes
Full-sample controls	No	Yes	Yes	No	Yes	Yes
Restricted sample controls	No	No	Yes	No	No	Yes
R-Square	0.09	0.10	0.46	0.09	0.10	0.49
Observations	33338	33338	4931	33338	33338	4931
Clusters	92	92	92	92	92	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 to 3 are instrumental variable regressions for the raw Macron vote share ; Columns 4 to 6 are instrumental variable regressions for the PS_j or "polling shift" dependent variable. Both use 1970-determined mean VLC in the municipality to instrument for the initial TH burden. Columns 1 and 4 include house values interacted with current TH rates, as well as other taxes depending on the VLC as controls. Columns 2 and 5 also include mean and median income, as well as population, as controls. Columns 3 and 6, in addition, include all deciles of income and the share of disposable income coming from pensions, capital income, and social transfers.

Table 3.B.5: Vote-stealing effects of the TH repeal promise

	(1) (Le Pen)	(2) (Fillon)	(3) (Melenchon)	(4) (Abstention)
Initial TH burden(EUR '000s)	-4.16****	0.13	0.00	1.68**
	(0.79)	(0.69)	(0.36)	(0.84)
Home value controls	Yes	Yes	Yes	Yes
Local tax controls	Yes	Yes	Yes	Yes
Full sample controls	Yes	Yes	Yes	Yes
R-Square	0.02	0.05	0.02	0.03
Observations	33338	33338	33338	33338
Clusters	92	92	92	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 to 4 are instrumental variable regressions for the PE_j or "prediction error" dependent variable, where the dependent variable is the excess vote share of each of the three main contenders besides E. Macron, and excess abstention over the prediction. All use 1970-determined mean VLC in the municipality to instrument for the initial TH burden. All specifications include house values interacted with current TH rates, other taxes depending on the VLC, as well as mean income per capita, and local population, as controls.

Table 3.B.6: Impact of the TH on Macron vote share: the role of home-ownership

	(OLS)	(IV)	(IV)	(OLS)	(IV)	(IV)
	Raw Macron vote share			"Polling shift" excess vote share		
Initial TH burden (EUR '000s) Initial TH burden	0.21	8.11*	4.34	1.36	6.28**	3.58
	(2.48)	(4.88)	(4.89)	(1.57)	(2.92)	(2.64)
Home ownership x Initial TH burden	9.80***	0.64	3.15	5.58***	-0.05	1.17
	(3.24)	(6.13)	(6.27)	(1.92)	(3.52)	(3.22)
Home value controls	No	Yes	Yes	No	Yes	Yes
Local tax controls	No	Yes	Yes	No	Yes	Yes
Full-sample controls	No	No	Yes	No	No	Yes
R-Square	0.07	0.10	0.11	0.06	0.09	0.10
Observations	35194	33335	33335	35194	33335	33335
Clusters	96	92	92	96	92	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 to 3 use the raw Macron vote share ; Columns 4 to 6 use PS_j or "polling shift" dependent variable. Columns 1 and 4 are OLS regressions; columns 2 to 3 and 5 to 6 are instrumental variable regressions, using 1970-determined mean VLC in the municipality to instrument for the initial TH burden and the interaction of the mean VLC with the home ownership rate to instrument for the interaction term. All specifications include controls for home ownership rates, house values interacted with current TH rates, other taxes depending on the VLC, as well as mean and median income, and local population, as controls.

Table 3.B.7: Impact of the TH on Macron vote share: the role of exemptions

	(OLS)	(OLS)	(IV)	(IV)
Initial TH burden	7.21****	5.27****	7.50****	5.12****
	(1.04)	(0.60)	(0.65)	(0.88)
Exempt share x TH burden	-0.22**	-0.14***	-0.21****	-0.10
	(0.09)	(0.04)	(0.06)	(0.06)
Home value controls	No	No	Yes	Yes
Local tax controls	No	No	Yes	Yes
Income controls	Yes	Yes	No	Yes
R-Square	0.09	0.08	0.09	0.10
Observations	35197	35197	33338	33338
Clusters	96	96	92	92

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 uses the raw Macron vote share ; Columns 2 to 4 use the PS_j or "polling shift"-corrected dependent variable. Columns 1 and 2 are OLS regressions; columns 3 and 4 are instrumental variable regressions, using 1970-determined mean VLC in the municipality to instrument for the respective measure of the TH burden and the interaction of the mean VLC with the exempted share to instrument for the interaction term. All specifications except column 3 include controls for population and mean income. Columns 3 and 4 also include as controls house values interacted with current TH rates, other taxes depending on the VLC.

Table 3.B.8: Impact of the TH on Macron vote share: the role of town size

	(OLS)	(OLS)	(OLS)	(IV)	(IV)
Initial TH burden	5.21****	3.79****			4.13****
	(1.00)	(0.49)			(0.73)
Pop. ('000) x TH burden	-0.02**	-0.01			0.12
	(0.01)	(0.01)			(0.09)
TH/income ratio			0.59****	0.95****	
			(0.11)	(0.19)	
Pop ('000) x TH/Income			-0.00	-0.04	
			(0.00)	(0.04)	
Home value controls	No	No	No	Yes	Yes
Local tax controls	No	No	No	Yes	Yes
Income controls	Yes	Yes	Yes	Yes	Yes
R-Square	0.08	0.08	0.11	0.13	0.10
Observations	35197	35197	31639	30068	33338
Clusters	96	96	96	92	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 uses the raw Macron vote share ; Columns 2 to 5 use the PS_j or "polling shift" dependent variable. Columns 1 to 3 are OLS regressions; columns 4 and 5 are instrumental variable regressions, using 1970-determined mean VLC in the municipality to instrument for the respective measure of the TH burden and the interaction of the mean VLC with population size to instrument for the interaction term. All specifications include controls for population and mean income. Columns 4 and 5 include as controls house values interacted with current TH rates, other taxes depending on the VLC.

Table 3.B.9: Impact of the TH, with province Fixed effects

	(OLS)	(IV)	(IV)	(OLS)	(IV)	(IV)
	Raw Macron vote share			"Polling shift" excess vote share		
Initial TH burden	8.261****	9.335****	5.626****	5.548****	5.833****	3.630****
	(0.581)	(0.728)	(0.770)	(0.389)	(0.538)	(0.534)
Province FE	Yes	Yes	Yes	Yes	Yes	Yes
Home value controls	No	Yes	Yes	No	Yes	Yes
Local tax controls	No	Yes	Yes	No	Yes	Yes
Full-sample controls	No	No	Yes	No	No	Yes
R-Square	0.259	0.102	0.138	0.132	0.0656	0.0828
Observations	35196	33338	33338	35196	33338	33338
Clusters	95	92	92	95	92	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. All columns include province fixed-effects to account for spatial correlation in 2012-2017 electoral shifts that may be correlated to province-level shocks to the initial housing tax burden or the instrument. Columns 1 to 3 use the raw Macron vote share; Columns 4 to 6 use PS_j, the "polling shift"-corrected dependent variable. Columns 1 and 4 are OLS regressions; columns 2, 3, 5 and 6 are instrumental variable regressions, using 1970-determined mean VLC in the municipality to instrument for the initial TH burden. All specifications include province fixed-effects. Columns 2 to 3 and 5 to 6 include controls for house values interacted with current TH rates, as well as other taxes depending on the VLC. Columns 3 and 6 include mean income and local population as controls.

Table 3.B.10: Impact of the TH in spatial first differences: comparing contiguous towns

	(OLS)	(OLS)	(IV)	(OLS)	(OLS)	(IV)
	Raw Macron vote share			"Polling shift" excess vote share		
SFD(Initial TH burden)	2.181****	2.654****	6.483****	1.773****	2.148****	4.195****
	(0.293)	(0.262)	(0.495)	(0.201)	(0.191)	(0.349)
Home value controls	No	No	Yes	No	No	Yes
Local tax controls	No	No	Yes	No	No	Yes
Income controls	No	Yes	No	No	Yes	No
R-Square	0.00445	0.0874	0.0147	0.00359	0.0389	0.00724
Observations	35180	31614	33321	35180	31614	33321
Clusters	96	96	92	96	96	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. Columns 1 to 3 use the raw Macron vote share; Columns 4 to 6 use PS_j or "polling shift" dependent variable. Columns 1, 2, 3 and 5 are OLS regressions; columns 4 and 6 are instrumental variable regressions, using 1970-determined (spatial first difference of) mean VLC in the municipalities to instrument for the initial (spatial first difference) TH burden. All variables, dependent and independent, correspond to $\Delta Y_j = Y_j - \bar{Y}_{N(j)}$, with $\bar{Y}_{N(j)}$ the mean value of Y in the set of municipalities contiguous to j, weighted by population. All IV specifications include controls for (the spatial first-difference of) house values interacted with current TH rates, and other taxes depending on the VLC. Columns 3 and 6 include spatial first differences of mean and median income, and local population, as controls.

Table 3.B.11: Alternative measures of the TH burden

	(1)	(2)	(3)	(4)	(5)
TH receipts/inhab.	2.643^{****}				
	(0.341)				
TH receipts/regist.		2.512^{****}			
		(0.295)			
TH/mean income			2.531^{****}		
			(0.454)		
TH/median income				2.031^{****}	
				(0.300)	
Government measure					1.396^{****}
					(0.142)
Home value controls	Yes	Yes	Yes	Yes	Yes
Local tax controls	Yes	Yes	Yes	Yes	Yes
Observations	33338	33338	30068	30068	28264
Clusters	92	92	92	92	92

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. All columns are instrumental variable regressions, using 1970-determined mean VLC in the municipality to instrument for the initial TH burden. All specifications include controls for house values interacted with current TH rates, and other taxes depending on the VLC.

Table 3.B.12: Placebo analysis: impact in 2012 election

	(Melenchon)	(Hollande)	(Bayrou)	(Sarkozy)	(Le Pen)
Initial TH burden	-0.19	0.48	-0.77***	-0.48	-0.91*
	(0.31)	(0.51)	(0.28)	(0.47)	(0.46)
Full sample controls	Yes	Yes	Yes	Yes	Yes
R-Square	0.00	0.00	0.00	0.01	0.01
Observations	32878	32878	32878	32878	32878
Clusters	96	96	96	96	96

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, **** $p < 0.001$

Standard errors in parentheses, clustered at the province level. All columns are OLS regressions, and include controls for population size and median income in 2012. The dependent variable is $PE_{j,k,2012}$, the prediction error between a candidate's actual 2012 vote share and its expected vote share based on a regression of 2012 results on 2007 presidential election vote shares for all candidates in the town. The independent variable of interest is the average housing tax burden per household as of 2011.

3.C Data appendix for Chapter 3

I use data from a variety of administrative sources in France to construct my outcome variables and key sources of variation at the municipality-level in the initial housing tax burden.

Units of observation The unit of observation is a "2017 municipality" (i.e. a municipality for which there exists available 2017 electoral data). In 2017, all 35,281 municipalities in France's mainland territory reported exhaustive electoral data.[34] Municipalities in France have been subject to numerous and continuous mergers and modifications. To aggregate municipalities at a consistent level, I use data from INSEE's *Code officiel géographique* files that include all mergers between municipalities that occurred from 2015 to 2017 (recorded in response to a law facilitating inter-municipal mergers), and manually add mergers that occurred from 2012 to 2015. This allows me to aggregate past 2012 electoral results (from municipalities existing in May 2012), and housing tax burdens known at the time of the election (from municipalities existing as of January 1st, 2016), to the level of 2017 municipalities, matching 99.9 percent of the electorate of 2017 mainland municipalities (35197 municipalities) to available data on past electoral data, housing tax receipts per household, average assessed values, town population, and average income.

Due to statistical secrecy constraints, data are not disclosed by INSEE, the French national statistical institute, when areas or subgroups with less than fifty households or a hundred individuals can be re-identified. Thus, while the baseline regression sample includes 35197 municipalities with available data, the smallest 11 percent of municipalities (around 3700) drop from the sample when adding controls for median income; and only the largest c. 5000 municipalities remain when adding controls for income distribution by deciles and composition of income by sources.

Electoral outcomes My main outcome variable is the share of registered voters who supported Emmanuel Macron in the first round of the French presidential election in April 2017. Exhaustive data on electoral outcomes for all municipalities in France in the first and second

[34]I exclude the cases of French citizens in foreign countries and French territories outside mainland France (except Corsica), since French citizens abroad are exempt from the housing tax, and specific exemptions and deductions are granted to French territories outside the mainland.

rounds of the 2017 and 2012 presidential elections are gathered by the French Interior Ministry, using figures from the national electoral commission (*Commission nationale de recensement des votes*). Similarly, I use data on 2012 first-round electoral results available from the Interior Ministry to construct two alternative predicted vote shares for Emmanuel Macron in 2017: one based on expected shifts from 2012 candidates to Emmanuel Macron inferred from opinion polls; and one based on linearly projecting Macron's electoral support on the vote share of all candidates in the 2012 election. Out of all municipalities that report 2017 electoral data, more than 80 percent only include one polling place; 60 percent of registered voters live in municipalities with less than 10 individual polling places.[35]

Local taxation Local taxation data are available for the year 2016 from the *Recensement des Elements d'Imposition a la Fiscalite Locale* (REI) database, designed by the French Treasury. I obtain detailed information on *Taxe d'habitation* tax rates, tax revenues, average assessed rental values (VLC or *Valeur Locative Cadastrale*), as well as municipality-specific deduction rates, and the number of beneficiaries of all deductions and exemptions. I use data from 2016, since the housing tax is paid in November each year, and the 2016 tax bill was the last known to households at the time of the April 2017 presidential election. The data include such details for all local taxes, including the TH, but also the *taxe fonciere* (an individual property tax) and the *contribution economique territoriale* and its successor, the *contribution sur la valeur ajoutee des entreprises*, a local corporate business tax. Since these other levies depend on the *Valeur Locative Cadastrale*, I control for receipts from these taxes when using the VLC as an instrument, to satisfy the exclusion restriction that the VLC does not affect voter support through e.g. other taxes that might have been expected to be changed by the election (although no candidate planned to reform any of the other local taxes based on the VLC).
My preferred measure for the average value of the housing tax burden per household in 2016, TH_j^1 divides the sum of TH receipts (at the municipal and inter-municipal level) by the number of households subject to the tax in a municipality. I include the collection fee F_t imposed by the national government, as well as a few small specific "historical" taxes as-

[35]Precincts or polling places, which count 670 registered voters on average and are recommended not to exceed a maximum of c. 1,000 voters, are the smallest unit level at which votes are recorded, but do not correspond to any administrative tax unit. Votes are also recorded at the municipality level, which is the relevant housing tax administrative layer. The average municipality counts 1,260 registered voters.

sociated with the TH, levied on the same tax return using the same base, which were also subject to the Macron measure.[36] As shown in figure 3.A.2, my main measure lines up almost exactly with an *ex post* estimate of the average savings per household, provided by a government document published in 2018, after the Macron election, quantifying the benefit of the housing tax cut at the municipality level.[37]

Housing market values There does not exist a publicly available set of current market-rate housing values at the municipality level in France. To construct conditionally exogenous variation and control for the role of current housing values in household sorting and local ideological alignment, as described in detail in subsection 3.4.1, I use the *Demande de Valeurs Foncieres* (DVF) database provided by the French Finance ministry.[38] The DVF database is an exhaustive repository of c. five million housing transactions that intervened in France over a period of five years (2014 to 2019), recording their exact location, total value, square footage, and, for some observations, specific characteristics of the sale. I match geo-coded transactions to their municipality, and use the full database[39] to compute median and mean home values, and price per square meter, for home sales in each municipality over the five-year window centered around the 2016-2017 period of interest. The (log) VLC based on 1970 values, and the current (log) median home value line up relatively closely, with a correlation of 0.57 (0.52 when using (log) mean home values).

Municipality-level characteristics To control for composition effects at the municipality level and test for heterogeneous responses, I also include municipality-level socio-demographic

[36] These notably include the so-called "taxe GEMAPI" and residence taxes specific to the Greater Paris area, which together represent less than 0.08 percent of overall TH receipts in the sample.

[37] This alternative measure, which has a 95 percent correlation with my baseline TH_j^1 figure in the available subset of c. 30,000 municipalities, is not used as the main dependent variable, for two reasons. First, it was not known to households at the time of the election; second, it includes the impact of potential variation in local tax rates between 2016 and 2018. In the robustness section, however, I experiment with a series of alternative measures of the TH burden, including this *ex post* government estimate.

[38] See Casanova Enault, Boulay, and Coulon (2019) for a detailed description of the DVF database and potential caveats in the reporting of the data.

[39] I only include sales of houses and apartments, and exclude sales of unbuilt land and smaller units like garages. I also exclude sales of units below 10 square meters and winsorize the sample by excluding extreme values for the price per sq. meter.

from a variety of sources.

The primary source is the FiLoSoFi database (*Fichiers Localises Sociaux et Fiscaux* or *Localized Social and Fiscal Files*) in 2016. These data, combined by INSEE, the French statistical institute, are designed specifically to cover all taxpayers subject to the TH. They include inputs from several sources, notably tax files, pension schemes, and social security services, as well as the full population census. The FiLoSoFi files provide detailed information on population composition, age structure, income levels by decile, the share of disposable income received from capital, labor, or pensions, as well as the share of income paid in taxes or received in welfare benefits. When including additional controls, the sample size varies slightly according to data availability, because INSEE does not disclose municipality-level data when it can be used to re-identify information about less than fifty households or a hundred individuals.

The Filosofi database provide median household income in a municipality. I also obtain mean income per capita at the municipality level, as well as local population counts, from a government dataset specifying criteria for the attribution of inter-governmental grants, the *Dotation Globale des Collectivites Locales* dataset.

I construct a data file of aggregate income tax payments per municipality from the *IRCOM* (*Impot sur le Revenu des Communes*) database, provided for each province, by the French Direction Generale des Finances Publiques.

I use Housing Census data on municipalities from the 2017 Supplemental Housing Survey to obtain the share of homeowners, the share of renters, and the share of secondary residences in a municipality. Finally, I use geographic adjacency files from OpenStreetMap, a collaborative open-source geographic information systems project providing the identifiers of all neighbouring municipalities for each municipality in my sample, enabling me to identify and compare neighbouring municipalities in a spatial first differences design.

Nationwide sources : To construct a pre-reform prediction for the Macron vote share at the local level, I rely on pre-announcement polling data from an Elabe poll published three days before the TH tax cut announcement, including a decomposition of voting intentions depending on the respondent's 2012 vote.[40] When demonstrating the importance of the TH

[40] Intentions de vote – Election présidentielle 2017, *Elabe etudes et sondages*, February 21st, 2017.

effect for Macron's support at the aggregate level, I avail myself of a variety of additional sources. In particular, I use Internet search data from Google Trends; polling data from IFOP's rolling daily poll; and prediction market data from the website PredictIt, when running event studies on estimated voter interest, victory chances, and polling results around the day when Macron announced the reform in February 2017.

3.D Details of the housing tax computation

The housing tax is part of four main sources of local tax revenue for municipalities in France. Breuillé, Duran-Vigneron, and Samson (2018) provide a detailed overview of the various components of municipal taxation in France. Municipalities also rely on direct government grants to fund their operations and investment projects. As of 2017, local tax revenue for municipalities stood at EUR 52 billion out of total receipts of c.EUR 90 billion, with government grants representing most of the remainder. The TH is collected on residents of furnished units. Residents of occupied units are defined as the owner-occupier or the tenant occupying the housing unit as of January 1st of each year; for vacant units and second homes, the tax is paid by owners. In 2018, the French government made the source code for the computation available as part of its open data policy, at the following address.

Tax base

The tax base of the TH is an assessed annual renting value of the unit, the *valeur locative cadastrale* (VLC). The VLC assessment, used for the determination of all local taxes in France, such as the individual property taxes on developed and undeveloped land is the product of an estimated weighted square footage, and an imputed rental rate per square meter.

Weighted square footage The weighted square footage (*surface ponderee*) computation started by measuring the actual square footage of a housing unit in 1970 (or at the time of construction). Each housing unit was then ranked into one a "quality category", corresponding to one of eight coarse levels on a ladder from "insalubrious" to "luxury", as classified arbitrarily in 1970 by civil servants in each town. The assessors then re-weighted the measured footage, using a formula over-weighting the first 20 square meters and under-weighting additional square footage beyond a threshold that varies according to the unit's quality, to account for category-specific estimated "decreasing returns" in rental services. The "weighted square footage" thus computed was adjusted upwards for the presence of what were deemed in 1970 to be "comfort" elements, like the presence of a bathtub (5 additional sq. m or 50 sq. ft), sinks (3 sq. m), running water (10 sq. m), or electricity (2 sq. m), each adding a fixed number of "weighted square footage".

The resulting weighted area is adjusted downwards or upwards by 5 percent in the presence of an elevator in 1970, and by two multiplicative factors (up to minus or plus 20 percent each): the "maintenance coefficient", which accounts for the age, maintenance requirements and overall condition of the unit – as of 1970; and the "peculiar situation coefficient", which takes into account its relative location within a municipality. Both adjustments were made with respect to the condition of the unit and the desirability of its location according to municipal employees as of January 1st, 1970.

Rental rates The weighted square footage obtained after all previous steps was then multiplied by a municipality- and category-specific "rental rate", defined in 1970, to obtain the imputed rental value. In each town, a few "representative units" (5.2 percent of units nationally) were assessed for each category, with municipal assessors recording either their actual market rental rate, observed as of January 1st, 1970, or an imputed rate (a constant return on the unit's last sale price as of 1970) for vacant units that were not currently rented and for owner-occupied units. Each non-reference housing unit was then allocated the municipality- and category-specific rental rate (per square meter of weighted square footage) of the reference unit. Multiplying the "rental rate" by the "imputed square footage" of the unit yielded a nominal VLC as of January 1970, or *VLC70*. Newly built housing after 1970 is classified into one of the eight quality categories, and then attributed a virtual *VLC70* according to the exact same process.

Revisions of the VLC - or absence thereof Officially, a July, 1, 1974 law calls for the revision of the *VLC70* every three years, as well as changes to the imputed square footage of a unit in the case of remodellings or additions. However, the VLC70 were never revised. They were "actualized" once in 1980, using a common province-level multiplicative adjustment factor, yielding somewhat updated VLCs, nicknamed *VLC80*.[41]

Together, the Senate, National Assembly, and two advisory and control bodies, the *Cour des*

[41] There are 96 provinces or *departements* of varying size and population in metropolitan France, each of which contains more than 350 municipalities on average. The maximum decadal nominal actualization coefficient applied in 1980 was 85 percent, in Paris. The minimum was 41 percent, in the Gers. The standard deviation of nominal decadal actualization across provinces was only eight percent. As a point of comparison, cumulative CPI inflation was over 170 percent over the same 1970-1980 period, according to INSEE, the French national statistical institute.

Comptes, and *Conseil des Prelevements Obligatoires*, authored more than twenty reports dedicated to the need for a general revision of the imputed rental values from 1990 to 2017. A July 1990 bill[42] was the only attempt at wholly revising rental values. Updated market rate assessments were collected throughout the French territory, leading to the existence of a shadow rental value database, the *VLC90*, at the French revenue service (then called the *Direction Generale des Impots*). According to the planned revision, rental values would have been updated upwards by more than 50 percent for at least 7 percent of units, and higher-income households would have been hit substantially harder.[43] However, after years of fierce debates, the government abandoned the project in the late 1990s, in view of the substantial implied redistributive effects. A small scale experiment to revise some VLC in five out of ninety-five provinces was also conducted in 2011, but did not lead to an updating of VLCs used in the housing tax.

The absence of local revisions of rental rates and imputed square footage over more than 45 years led to a substantial geographic disconnect between current market valuations and the formulaic assessments. Even the "weighted square footage" had rarely been updated as of 2017, even though virtually all city center housing units had undergone substantial remodellings and additions of "comfort" elements, such as elevators, electricity, bathrooms, or running water, over the previous five decades. The *VLC* are sometimes adjusted on a case-by-case basis in the presence of egregious mistakes, or changes in the rental value due to substantial architectural changes, although there is no penalty for failing to disclose a remodelling addition to the tax administration. Indeed, smaller municipalities often only receive information from the national revenue service (DGFIP) about the *aggregate* VLC in the town.

Tax rates

Tax rates The municipal tax rate applied to the base, on the other hand, is defined by a vote in city councils every year. As of 2016, there were two main levels of the tax in force,

[42] *Loi n.90-669 du 30 juillet 1990 relative à la révision générale des évaluations des immeubles retenus pour la détermination des bases des impôts directs locaux*

[43] See the June 2012 Senate Ways and Means committee report on revising the VLC, Rapport d'information fait au nom de la Commission des Finances sur la révision des valeurs locatives professionnelles et commerciales

at the municipality or *commune* level, and at the level of the "inter-municipality" or *EPCI* (a syndicate of towns who jointly provide a variety of local public services). Additional tax rates are imposed in some municipalities to fund flooding prevention or other special services.

Applicable adjustments and deductions The tax base, obtained by multiplying the rental value by the imputed square footage, is adjusted for a number of deductions, most of which correspond to reductions for the number of dependents living in the household, or special exemptions for older citizens, lower-income households, or people with disabilities. Most deductions are aimed at reducing the burden of the tax for larger and/or underprivileged households. They are defined as a fixed deduction rate (e.g. 10% per additional child) multiplied by the *average* VLC in the municipality, $\overline{\text{VLC}}_{jt}$.[44] Specific categories of households (people above 60, disabled individuals, and widowers) with an income below a threshold close to the national poverty line, adjusted for household size, are exempt from the tax. Finally, the tax is subject to a ceiling of 3.44% of the household's annual fiscal income. Most deductions and ceilings are "paid for" by the central government, through inter-governmental grants to municipal authorities compensating for the corresponding lost revenue.[45] Since only 43 percent of households paid any progressive income tax (*Impot sur le revenu*) in 2016, the TH was also the only broadly applicable and salient tax remitted by households at almost all income levels above the poverty line.[46]

[44]One national mandatory deduction is increasing in the number of dependents: the deduction rate can be adapted by municipalities from 10 to 20 percent for the first two children, and from 15 to 25 percent per child after the third. Two deductions, up to a ceiling of fifteen percent each, can be voted upon locally in city councils: one applies to all households, and the other to households below a means-tested threshold.

[45]In 2017, out of EUR 23.9bn (slightly less than two percent of aggregate disposable income) of total TH receipts, EUR 1.7bn corresponded to the inter-governmental compensation for exempted households, and EUR 4bn to the inter-governmental compensation of deductions, leaving 18.2bn (76 percent) directly paid by households locally.

[46]Most households not subject to the progressive income tax but receiving wage or pension income pay a c. 20 percent Value Added Tax on consumption, a c. 10 percent flat wage tax (*CSG/CRDS*), and social security contributions covering health and unemployment insurance ranging from 30 to 60 percent of total labor costs. Nonetheless, because they are either directly subtracted from gross income or included in sale prices, these levies are generally not as salient as direct tax payments from the point of view of households (see e.g. Bozio, Breda, and Grenet (2017)).

Bibliography

Abdel-Rahman, Hesham and Masahisa Fujita (1990). "Product variety, Marshallian externalities, and city sizes". *Journal of Regional Science* 30.2, pp. 165–183.

Abdel-Rahman, Hesham M (1994). "Economies of scope in intermediate goods and a system of cities". *Regional Science and Urban Economics* 24.4, pp. 497–524.

Abdel-Rahman, Hesham M and Masahisa Fujita (1993). "Specialization and diversification in a system of cities". *Journal of Urban Economics* 33.2, pp. 189–222.

Ades, Alberto F and Edward L Glaeser (1995). "Trade and circuses: explaining urban giants". *The Quarterly Journal of Economics* 110.1, pp. 195–227.

Aghion, Philippe and Patrick Bolton (1997). "A theory of trickle-down growth and development". *The Review of Economic Studies* 64.2, pp. 151–172.

Ahlfeldt, Gabriel M and Wolfgang Maennig (2015). "Homevoters vs. leasevoters: A spatial analysis of airport effects". *Journal of Urban Economics* 87, pp. 85–99.

Aidt, Toke S and Graham Mooney (2014). "Voting suffrage and the political budget cycle: Evidence from the London Metropolitan Boroughs 1902–1937". *Journal of Public Economics* 112, pp. 53–71.

Alesina, Alberto and Matteo Paradisi (2017). "Political budget cycles: Evidence from Italian cities". *Economics & Politics* 29.2, pp. 157–177.

Alpino, Matteo (2018). *On the electoral effectiveness of pre-election policy promises*. Tech. rep.

Aragonès, Enriqueta, Andrew Postlewaite, and Thomas Palfrey (2007). "Political reputations and campaign promises". *Journal of the European Economic Association* 5.4, pp. 846–884.

Arefeva, Alina et al. (2021). *Job Growth from Opportunity Zones*. Tech. rep.

Atkeson, Andrew and Ariel Burstein (2008). "Pricing-to-market, trade costs, and international relative prices". *American Economic Review* 98.5, pp. 1998–2031.

Autor, David, David Dorn, and Gordon H Hanson (2013). "The China syndrome: Local labor market effects of import competition in the United States". *American Economic Review* 103.6, pp. 2121–68.

Bagues, Manuel and Berta Esteve-Volart (2016). "Politicians' luck of the draw: Evidence from the Spanish Christmas lottery". *Journal of Political Economy* 124.5, pp. 1269–1294.

Bailey, Michael et al. (2018). "Social connectedness: Measurement, determinants, and effects". *Journal of Economic Perspectives* 32.3, pp. 259–80.

Bakker, Jan David (2018). "International trade and regional inequality". *CEPREMAP Working Paper* 1806.

Baum-Snow, Nathaniel and Justin Marion (2009). "The effects of low income housing tax credit developments on neighborhoods". *Journal of Public Economics* 93.5-6, pp. 654–666.

Berger, David, Nicholas Turner, and Eric Zwick (2020). "Stimulating housing markets". *The Journal of Finance* 75.1, pp. 277–321.

Berger, David W, Kyle F Herkenhoff, and Simon Mongey (forthcoming). "Labor market power". *American Economic Review*.

Bernard, Andrew B, Raymond Robertson, and Peter K Schott (2010). "Is Mexico a lumpy country?" *Review of International Economics* 18.5, pp. 937–950.

Best, Michael Carlos and Henrik Jacobsen Kleven (2018). "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK". *The Review of Economic Studies* 85.1, pp. 157–193.

Blouri, Yashar, Simon Büchler, and Olivier Schöni (2021). "The geography of housing subsidies". *MIT Center for Real Estate Research Paper* 21/05.

Bombardini, Matilde, Giovanni Gallipoli, and Germán Pupato (2012). "Skill dispersion and trade flows". *American Economic Review* 102.5, pp. 2327–48.

Bono, Pierre-Henri and Alain Trannoy (2019). "The Impact of the 'Scellier'Income Tax Relief on Building Land Prices in France". *Economie et Statistique/Economics and Statistics* 507-508, pp. 91–114.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2021). *Revisiting Event Study Designs: Robust and Efficient Estimation*. Tech. rep.

Bozio, Antoine, Thomas Breda, and Julien Grenet (2017). "Incidence of social security contributions: evidence from France". *Paris School of Economics Working Paper*.

Bozio, Antoine et al. (2015). *Réformer les aides personnelles au logement*. Institut des politiques publiques (IPP).

Brakman, Steven and Charles Van Marrewijk (2013). "Lumpy countries, urbanization, and trade". *Journal of International Economics* 89.1, pp. 252–261.

Brender, Adi and Allan Drazen (2008). "How do budget deficits and economic growth affect reelection prospects? Evidence from a large panel of countries". *American Economic Review* 98.5, pp. 2203–20.

Breuillé, Marie-Laure, Pascale Duran-Vigneron, and Anne-Laure Samson (2018). "Inter-municipal cooperation and local taxation". *Journal of Urban Economics* 107, pp. 47–64.

Brewer, Mike et al. (2019). "The curious incidence of rent subsidies: Evidence of heterogeneity from administrative data". *Journal of Urban Economics* 114, pp. 103–198.

Brunner, Eric, Stephen L Ross, and Ebonya Washington (2011). "Economics and policy preferences: causal evidence of the impact of economic conditions on support for redistribution and other ballot proposals". *Review of Economics and Statistics* 93.3, pp. 888–906.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro (2015). *The impact of trade on labor market dynamics*. Tech. rep. National Bureau of Economic Research.

Campbell, John Y and Robert J Shiller (1988). "The dividend-price ratio and expectations of future dividends and discount factors". *The Review of Financial Studies* 1.3, pp. 195–228.

Campbell, Sean D et al. (2009). "What moves housing markets: A variance decomposition of the rent–price ratio". *Journal of Urban Economics* 66.2, pp. 90–102.

Carozzi, Felipe, Christian Hilber, and Xiaolun Yu (2019). "The economic impacts of help to buy".

Casaburi, Lorenzo and Ugo Troiano (2016). "Ghost-house busters: The electoral response to a large anti–tax evasion program". *The Quarterly Journal of Economics* 131.1, pp. 273–314.

Casanova Enault, Laure, Guilhem Boulay, and Mathieu Coulon (2019). "Une aubaine pour les géographes? Intérêts des fichiers open DVF sur les transactions foncières et immobilières et précautions d'usage". *Cybergeo: European Journal of Geography*.

Caselli, Francesco and James Feyrer (2007). "The marginal product of capital". *The Quarterly Journal of Economics* 122.2, pp. 535–568.

Chandler, Tertius (1987). *Four thousand years of urban growth: an historical census.* Lewiston New York St. Davids University Press 1987.

Charpentier, Arthur (2017). "Les marchés prédictifs comme technique de prévision". *Risques* 111.

Chen, Henry et al. (2010). "Buy local? The geography of venture capital". *Journal of Urban Economics* 67.1, pp. 90–102.

Chen, Jiafeng, Edward L Glaeser, and David Wessel (2019). *The (non-) effect of opportunity zones on housing prices*. National Bureau of Economic Research.

Chinco, Alex and Christopher Mayer (2016). "Misinformed speculators and mispricing in the housing market". *The Review of Financial Studies* 29.2, pp. 486–522.

Chodorow-Reich, Gabriel, Plamen T Nenov, and Alp Simsek (2021). "Stock market wealth and the real economy: A local labor market approach". *American Economic Review* 111.5, pp. 1613–57.

Chor, Davin (2010). "Unpacking sources of comparative advantage: A quantitative approach". *Journal of International Economics* 82.2, pp. 152–167.

Ciccone, Antonio and Robert E Hall (1996). "Productivity and the Density of Economic Activity". *American Economic Review* 86.1, pp. 54–70.

Cingano, Federico and Paolo Pinotti (2016). "Trust, firm organization, and the pattern of comparative advantage". *Journal of International Economics* 100, pp. 1–13.

Collinson, Robert and Peter Ganong (2018). "How do changes in housing voucher design affect rent and neighborhood quality?" *American Economic Journal: Economic Policy* 10.2, pp. 62–89.

Coşar, A Kerem and Pablo D Fajgelbaum (2016). "Internal geography, international trade, and regional specialization". *American Economic Journal: Microeconomics* 8.1, pp. 24–56.

Costinot, Arnaud (2009). "On the origins of comparative advantage". *Journal of International Economics* 77.2, pp. 255–264.

Costinot, Arnaud, Dave Donaldson, and Ivana Komunjer (2011). "What goods do countries trade? A quantitative exploration of Ricardo's ideas". *The Review of Economic Studies* 79.2, pp. 581–608.

Coulomb, Renaud and Marc Sangnier (2014). "The impact of political majorities on firm value: Do electoral promises or friendship connections matter?" *Journal of Public Economics* 115, pp. 158–170.

Courant, Paul N and Alan V Deardorff (1992). "International Trade with Lumpy Countries". *Journal of Political Economy*, pp. 198–210.

— (1993). "Amenities, nontraded goods, and the trade of lumpy countries". *Journal of Urban Economics* 34.2, pp. 299–317.

Coval, Joshua D and Tobias J Moskowitz (1999). "Home bias at home: Local equity preference in domestic portfolios". *The Journal of Finance* 54.6, pp. 2045–2073.

Criscuolo, Chiara et al. (2019). "Some causal effects of an industrial policy". *American Economic Review* 109.1, pp. 48–85.

Cruz, Cesi et al. (2018). *Making policies matter: Voter responses to campaign promises*. Tech. rep. National Bureau of Economic Research.

Davis, Donald R and Jonathan I Dingel (2014). *The comparative advantage of cities*. Tech. rep. National Bureau of Economic Research.

De La O, Ana L (2013). "Do conditional cash transfers affect electoral behavior? Evidence from a randomized experiment in Mexico". *American Journal of Political Science* 57.1, pp. 1–14.

Debaere, Peter (2004). "Does lumpiness matter in an open economy?: Studying international economics with regional data". *Journal of international Economics* 64.2, pp. 485–501.

Desmet, Klaus and Esteban Rossi-Hansberg (2013). "Urban accounting and welfare". *American Economic Review* 103.6, pp. 2296–2327.

Diamond, Rebecca and Tim McQuade (2019). "Who wants affordable housing in their backyard? An equilibrium analysis of low-income property development". *Journal of Political Economy* 127.3, pp. 1063–1117.

Dix-Carneiro, Rafael and Brian K Kovak (2015). "Trade liberalization and the skill premium: A local labor markets approach". *American Economic Review* 105.5, pp. 551–57.

Donaldson, Dave and Richard Hornbeck (2016). "Railroads and American economic growth: a "market access" approach". *The Quarterly Journal of Economics* 131.2, pp. 799–858.

Downs, Anthony (1957). "An economic theory of political action in a democracy". *Journal of Political Economy* 65.2, pp. 135–150.

Druckenmiller, Hannah and Solomon Hsiang (2018). *Accounting for unobservable heterogeneity in cross section using spatial first differences*. Tech. rep. National Bureau of Economic Research.

Duranton, Gilles and Diego Puga (2001). "Nursery cities: Urban diversity, process innovation, and the life cycle of products". *American Economic Review* 91.5, pp. 1454–1477.

— (2004). "Micro-foundations of urban agglomeration economies". *Handbook of regional and urban economics*. Vol. 4. Elsevier, pp. 2063–2117.

Duranton, Gilles and Matthew A Turner (2018). "Urban form and driving: Evidence from US cities". *Journal of Urban Economics* 108, pp. 170–191.

Edgerton, Jesse (2011). *The effects of taxation on business investment: New evidence from used equipment*. Federal Reserve Board Working Paper.

Eichholtz, Piet et al. (Mar. 2021). "The Total Return and Risk to Residential Real Estate". *The Review of Financial Studies* 34.8, pp. 3608–3646.

Elinder, Mikael, Henrik Jordahl, and Panu Poutvaara (2015). "Promises, policies and pocketbook voting". *European Economic Review* 75, pp. 177–194.

Ellison, Glenn, Edward L Glaeser, and William R Kerr (2010). "What causes industry agglomeration? Evidence from coagglomeration patterns". *American Economic Review* 100.3, pp. 1195–1213.

Eriksen, Michael D and Stuart S Rosenthal (2010). "Crowd-out effects of place-based subsidized rental housing: New evidence from the LIHTC program". *Journal of Public Economics* 94.11-12, pp. 953–966.

Eriksen, Michael D and Amanda Ross (2015). "Housing vouchers and the price of rental housing". *American Economic Journal: Economic Policy* 7.3, pp. 154–76.

Fack, Gabrielle (2006). "Are housing benefits an effective way to redistribute income? Evidence from a natural experiment in France". *Labour Economics* 13.6, pp. 747–771.

Faggio, Giulia, Olmo Silva, and William C Strange (2017). "Heterogeneous agglomeration". *Review of Economics and Statistics* 99.1, pp. 80–94.

Fajgelbaum, Pablo D and Amit K Khandelwal (2016). "Measuring the unequal gains from trade". *The Quarterly Journal of Economics* 131.3, pp. 1113–1180.

Fajgelbaum, Pablo D et al. (2019). "State taxes and spatial misallocation". *The Review of Economic Studies* 86.1, pp. 333–376.

Fally, Thibault (2015). "Structural gravity and fixed effects". *Journal of International Economics* 97.1, pp. 76–85.

Fang, Yu and James W Jawitz (2018). "High-resolution reconstruction of the United States human population distribution, 1790 to 2010". *Scientific data* 5, p. 180067.

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh (2019). *Affordable housing and city welfare*. Tech. rep. National Bureau of Economic Research.

Favilukis, Jack and Stijn Van Nieuwerburgh (2017). "Out-of-town Home Buyers and City Welfare". *The Journal of Finance*.

Feenstra, Robert C (1986). "Functional equivalence between liquidity costs and the utility of money". *Journal of Monetary Economics* 17.2, pp. 271–291.

Feldman, Stanley (1982). "Economic self-interest and political behavior". *American Journal of Political Science*, pp. 446–466.

Feldstein, Martin and Charles Horioka (1980). "Domestic Saving and International Capital Flows". *The Economic Journal* 90.358, pp. 314–329.

Fetzer, Thiemo (2019). "Did austerity cause Brexit?" *American Economic Review* 109.11, pp. 3849–86.

Fischel, William A (2009). *The homevoter hypothesis*. Harvard University Press.

Foremny, Dirk and Nadine Riedel (2014). "Business taxes and the electoral cycle". *Journal of Public Economics* 115, pp. 48–61.

Frank, Thomas (2007). *What's the matter with Kansas?: How conservatives won the heart of America*. Metropolitan Books.

Freedman, Matthew, Shantanu Khanna, and David Neumark (2021). *The Impacts of Opportunity Zones on Zone Residents*. Tech. rep. National Bureau of Economic Research.

French, Kenneth R and James M Poterba (1991). "Investor diversification and international equity markets". *American Economic Review* 81.2, pp. 222–226.

French, Scott (2017). "Revealed comparative advantage: What is it good for?" *Journal of International Economics* 106, pp. 83–103.

Gabaix, Xavier (1999a). "Zipf's Law and the Growth of Cities". *American Economic Review* 89.2, pp. 129–132.

— (1999b). "Zipf's law for cities: an explanation". *The Quarterly journal of economics* 114.3, pp. 739–767.

Ganong, Peter and Daniel Shoag (2017). "Why has regional income convergence in the US declined?" *Journal of Urban Economics* 102, pp. 76–90.

Garbinti, Bertrand, Jonathan Goupille-Lebret, and Thomas Piketty (2021). "Accounting for wealth-inequality dynamics: Methods, estimates, and simulations for France". *Journal of the European Economic Association* 19.1, pp. 620–663.

Gaubert, Cecile (2018). "Firm sorting and agglomeration". *American Economic Review* 108.11, pp. 3117–53.

Gaubert, Cecile, Patrick M Kline, and Danny Yagan (2021). *Place-Based Redistribution*. Tech. rep. National Bureau of Economic Research.

Gete, Pedro and Michael Reher (2018). "Mortgage supply and housing rents". *The Review of Financial Studies* 31.12, pp. 4884–4911.

Giacoletti, Marco (Mar. 2021). "Idiosyncratic Risk in Housing Markets". *The Review of Financial Studies* 34.8, pp. 3695–3741.

Glaeser, Edward and Joseph Gyourko (2018). "The economic implications of housing supply". *Journal of Economic Perspectives* 32.1, pp. 3–30.

Glaeser, Edward L, Matthew E Kahn, and Jordan Rappaport (2008). "Why do the poor live in cities? The role of public transportation". *Journal of Urban Economics* 63.1, pp. 1–24.

Greenwald, Daniel L and Adam Guren (2019). *Do credit conditions move house prices?* Tech. rep. MIT.

Grossman, Gene M and Elhanan Helpman (n.d.). "Identity politics and trade policy". *Review of Economic Studies* Forthcoming ().

Gruber, Jonathan, Amalie Jensen, and Henrik Kleven (2021). "Do people respond to the mortgage interest deduction? Quasi-experimental evidence from Denmark". *American Economic Journal: Economic Policy* 13.2, pp. 273–303.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai (2013). "Superstar cities". *American Economic Journal: Economic Policy* 5.4, pp. 167–99.

Hall, Robert E and Dale W Jorgenson (1967). "Tax policy and investment behavior". *American Economic Review* 57.3, pp. 391–414.

Helpman, Elhanan (1998). "The size of regions". *Topics in public economics: Theoretical and applied analysis*, pp. 33–54.

Herreno, Juan (2020). *The Aggregate Effects of Bank Lending Cuts*. Tech. rep. Columbia University.

Hilber, Christian, Andreas Mense, et al. (2021). *Why have house prices risen so much more than rents in superstar cities?*

Hilber, Christian AL and Tracy M Turner (2014). "The mortgage interest deduction and its impact on homeownership decisions". *Review of Economics and Statistics* 96.4, pp. 618–637.

Holmes, Thomas J and John J Stevens (2004). “Spatial distribution of economic activities in North America”. *Handbook of Regional and Urban Economics*. Vol. 4. Elsevier, pp. 2797–2843.

Holmstrom, Bengt and Jean Tirole (1997). “Financial intermediation, loanable funds, and the real sector”. *the Quarterly Journal of economics* 112.3, pp. 663–691.

Hong, Harrison, Jeffrey D Kubik, and Jeremy C Stein (2008). “The only game in town: Stock-price consequences of local bias”. *Journal of Financial Economics* 90.1, pp. 20–37.

Hoopes, Jeffrey L, Daniel H Reck, and Joel Slemrod (2015). “Taxpayer search for information: Implications for rational attention”. *American Economic Journal: Economic Policy* 7.3, pp. 177–208.

House, Christopher L and Matthew D Shapiro (2008). “Temporary investment tax incentives: Theory with evidence from bonus depreciation”. *American Economic Review* 98.3, pp. 737–68.

Hsieh, Chang-Tai and Enrico Moretti (2019). “Housing constraints and spatial misallocation”. *American Economic Journal: Macroeconomics* 11.2, pp. 1–39.

Huet-Vaughn, Emiliano (2019). “Stimulating the vote: Arra road spending and vote share”. *American Economic Journal: Economic Policy* 11.1, pp. 292–316.

Inspection Generale des Finances (2019). *Évaluation du dispositif d’aide fiscale à l’investissement locatif Pinel*.

Ioannides, Yannis M and Henry G Overman (2003). “Zipf’s law for cities: an empirical examination”. *Regional science and urban economics* 33.2, pp. 127–137.

Irwin, Douglas A (2017). *Clashing Over Commerce: A History of US Trade Policy*. University of Chicago Press.

Jacobs, Jane (1969). *The Economics of Cities*. Random House.

Keesing, Donald B and Donald R Sherk (1971). “Population density in patterns of trade and development”. *American Economic Review* 61.5, pp. 956–961.

Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi (2015). "How do voters respond to information? Evidence from a randomized campaign". *American Economic Review* 105.1, pp. 322–53.

Kennedy, Patrick and Harrison Wheeler (2021). *Neighborhood-Level Investment from the US Opportunity Zone Program: Early Evidence*. Tech. rep.

King, Willford I (1927). "Wealth distribution in the continental United States at the close of 1921". *Journal of the American Statistical Association* 22.158, pp. 135–153.

Kline, Patrick and Enrico Moretti (2014a). "Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority". *The Quarterly journal of economics* 129.1, pp. 275–331.

— (2014b). "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs". *Annu. Rev. Econ.* 6.1, pp. 629–662.

Kotlikoff, Laurence J (1983). "National savings and economic policy: The efficacy of investment vs. savings incentives". *American Economic Review* 73.2, pp. 82–87.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen (2012). "The aggregate demand for treasury debt". *Journal of Political Economy* 120.2, pp. 233–267.

Krugman, Paul (1991). "Increasing returns and economic geography". *Journal of Political Economy* 99.3, pp. 483–499.

Krugman, Paul and Raul Livas Elizondo (1996). "Trade policy and the third world metropolis". *Journal of development economics* 49.1, pp. 137–150.

Kuchler, Theresa et al. (forthcoming). "Social proximity to capital: Implications for investors and firms". *Review of Financial Studies*.

Landau, Sarah Bradford and Carl W Condit (1999). *Rise of the New York skyscraper: 1865-1913*. Yale University Press.

Lanz, Rainer, Sebastien Miroudot, and Hildegunn K Nordås (2013). "Offshoring of tasks: Taylorism versus Toyotism". *The World Economy* 36.2, pp. 194–212.

Levitt, Steven D and James M Snyder Jr (1997). "The impact of federal spending on House election outcomes". *Journal of Political Economy* 105.1, pp. 30–53.

Levy, Antoine (2021). *The tax elasticity of capital gains: evidence from millions of sales in France*. Tech. rep. MIT.

Lewis-Beck, Michael S and Mary Stegmaier (2000). “Economic determinants of electoral outcomes”. *Annual Review of Political Science* 3.1, pp. 183–219.

Lindbeck, Assar and Jörgen W Weibull (1993). “A model of political equilibrium in a representative democracy”. *Journal of Public Economics* 51.2, pp. 195–209.

Liu, Zheng, Pengfei Wang, and Tao Zha (2019). *A theory of housing demand shocks*. Tech. rep. National Bureau of Economic Research.

Lizzeri, Alessandro and Nicola Persico (2001). “The provision of public goods under alternative electoral incentives”. *American Economic Review* 91.1, pp. 225–239.

Lucas, Robert E (1990). “Why doesn’t capital flow from rich to poor countries?” *American Economic Review* 80.2, pp. 92–96.

Malpezzi, Stephen and Kerry Vandell (2002). “Does the low-income housing tax credit increase the supply of housing?” *Journal of Housing Economics* 11.4, pp. 360–380.

Manacorda, Marco, Edward Miguel, and Andrea Vigorito (2011). “Government transfers and political support”. *American Economic Journal: Applied Economics* 3.3, pp. 1–28.

Marshall, Alfred (1890). *The Principles of Economics*. MacMillan.

Metcalf, Gabriel (2018). “Sand castles before the tide? Affordable housing in expensive cities”. *Journal of Economic Perspectives* 32.1, pp. 59–80.

Michaels, Guy, Ferdinand Rauch, and Stephen J Redding (2013). “Task Specialization in US Cities from 1880 to 2000”. *Journal of the European Economic Association*.

— (2019). “Task Specialization in US Cities from 1880 to 2000”. *Journal of the European Economic Association* 17.3, pp. 754–798.

Moll, Benjamin (2014). “Productivity losses from financial frictions: Can self-financing undo capital misallocation?” *American Economic Review* 104.10, pp. 3186–3221.

Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg (2018). “Commuting, migration, and local employment elasticities”. *American Economic Review* 108.12, pp. 3855–90.

Moretti, Enrico (2012). *The new geography of jobs*. Houghton Mifflin Harcourt.

Nakamura, Ryohei (1985). "Agglomeration economies in urban manufacturing industries: a case of Japanese cities". *Journal of Urban economics* 17.1, pp. 108–124.

Nunn, Nathan (2007). "Relationship-specificity, incomplete contracts, and the pattern of trade". *The Quarterly Journal of Economics* 122.2, pp. 569–600.

Ohrn, Eric (2018). "The Effect of corporate taxation on investment and financial policy: evidence from the DPAD". *American Economic Journal: Economic Policy* 10.2, pp. 272–301.

Olsen, Edgar O and Jeffrey E Zabel (2015). "US housing policy". *Handbook of regional and urban economics*. Vol. 5. Elsevier, pp. 887–986.

Oswald, Florian (2019). "The effect of homeownership on the option value of regional migration". *Quantitative Economics* 10.4, pp. 1453–1493.

Pellegrino, Bruno, Enrico Spolaore, and Romain Wacziarg (2021). *Barriers to Global Capital Allocation*. Tech. rep. National Bureau of Economic Research.

Persson, Torsten and Guido Tabellini (2016). *Political economics*. MIT press.

Pflüger, Michael and Takatoshi Tabuchi (2016). "Comparative Advantage and Agglomeration of Economic Activity". 10273.

Pierce, Justin R and Peter K Schott (2012). "A concordance between ten-digit US Harmonized System Codes and SIC/NAICS product classes and industries". *Journal of Economic and Social Measurement* 37.1, 2, pp. 61–96.

Porter, Michael (2003). "The economic performance of regions". *Regional Studies* 37.6-7, pp. 549–578.

Portes, Richard and Helene Rey (2005). "The determinants of cross-border equity flows". *Journal of International Economics* 65.2, pp. 269–296.

Portes, Richard, Hélène Rey, and Yonghyup Oh (2001). "Information and capital flows: The determinants of transactions in financial assets". *European economic review* 45.4-6, pp. 783–796.

Poterba, James M (1984). "Tax subsidies to owner-occupied housing: an asset-market approach". *The quarterly journal of economics* 99.4, pp. 729–752.

Ramondo, Natalia, Andrés Rodríguez-Clare, and Milagro Saborío-Rodríguez (2016). "Trade, domestic frictions, and scale effects". *American Economic Review* 106.10, pp. 3159–84.

Rauch, James E (1991). "Comparative Advantage, Geographic Advantage and the Volume of Trade". *The Economic Journal* 101.408, pp. 1230–1244.

Reba, Meredith, Femke Reitsma, and Karen C Seto (2016). "Spatializing 6,000 years of global urbanization from 3700 BC to AD 2000". *Scientific data* 3, p. 160034.

Redding, Stephen J (2016). "Goods trade, factor mobility and welfare". *Journal of International Economics* 101, pp. 148–167.

Redding, Stephen J and Esteban Rossi-Hansberg (2017). "Quantitative spatial economics". *Annual Review of Economics* 9, pp. 21–58.

Redding, Stephen J and David E Weinstein (2020). "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences". *The Quarterly Journal of Economics* 135.1, pp. 503–560.

Reher, Michael (2021). "Finance and the supply of housing quality". *Journal of Financial Economics*.

Ricci, Luca Antonio (1999). "Economic geography and comparative advantage:: Agglomeration versus specialization". *European Economic Review* 43.2, pp. 357–377.

Roback, Jennifer (1982). "Wages, rents, and the quality of life". *Journal of Political Economy* 90.6, pp. 1257–1278.

Rodden, Jonathan (2010). "The geographic distribution of political preferences". *Annual Review of Political Science* 13, pp. 321–340.

Rodrik, Dani (2020). *Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-wing Populism*. Tech. rep. National Bureau of Economic Research.

Romalis, John (2004). "Factor proportions and the structure of commodity trade". *American Economic Review* 94.1, pp. 67–97.

Rosen, Sherwin (1979). "Wage-based indexes of urban quality of life". *Current issues in urban economics*, pp. 74–104.

Rosenthal, Stuart S and Stephen L Ross (2015). "Change and persistence in the economic status of neighborhoods and cities". *Handbook of Regional and Urban Economics*. Vol. 5. Elsevier, pp. 1047–1120.

Rosenthal, Stuart S and William C Strange (2001). "The determinants of agglomeration". *Journal of urban economics* 50.2, pp. 191–229.

— (2004). "Evidence on the nature and sources of agglomeration economies". *Handbook of regional and urban economics*. Vol. 4. Elsevier, pp. 2119–2171.

— (2008). "The attenuation of human capital spillovers". *Journal of Urban Economics* 64.2, pp. 373–389.

Saiz, Albert (2010). "The geographic determinants of housing supply". *The Quarterly Journal of Economics* 125.3, pp. 1253–1296.

Sato, Kazuo (1976). "The ideal log-change index number". *The Review of Economics and Statistics*, pp. 223–228.

Schuberth, Christopher J (1968). *The geology of New York City and environs*. Published for the American Museum of Natural History [by] Natural History Press.

Silva, JMC Santos and Silvana Tenreyro (2006). "The log of gravity". *The Review of Economics and Statistics* 88.4, pp. 641–658.

Sinai, Todd and Nicholas S Souleles (2005). "Owner-occupied housing as a hedge against rent risk". *The Quarterly Journal of Economics* 120.2, pp. 763–789.

Sinai, Todd and Joel Waldfogel (2005). "Do low-income housing subsidies increase the occupied housing stock?" *Journal of Public Economics* 89.11-12, pp. 2137–2164.

Singh, Divya (2019). "Do property tax incentives for new construction spur gentrification? Evidence from New York City". *112th Annual Conference on Taxation*. National Tax Association.

Sivitanides, Petros et al. (2001). "The determinants of appraisal-based capitalization rates". *Real Estate Finance* 18.2, pp. 27–38.

Slattery, Cailin and Owen Zidar (2020). "Evaluating state and local business incentives". *Journal of Economic Perspectives* 34.2, pp. 90–118.

Soltas, Evan J (2020). “The Price of Inclusion: Evidence from Housing Developer Behavior”.

Sommer, Kamila and Paul Sullivan (2018). “Implications of US tax policy for house prices, rents, and homeownership”. *American Economic Review* 108.2, pp. 241–74.

Spenkuch, Jörg L and David Toniatti (2018). “Political advertising and election results”. *The Quarterly Journal of Economics* 133.4, pp. 1981–2036.

Stewart, Charles (1939). “Income Capitalization as a Method of Estimating the Distribution of Wealth By Size Groups”. *Studies in Income and Wealth, Volume 3*. NBER, pp. 95–146.

Suárez Serrato, Juan Carlos and Owen Zidar (2016). “Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms”. *American Economic Review* 106.9, pp. 2582–2624.

Taylor, Paul J et al. (2008). “The transportability of job information across countries”. *Personnel Psychology* 61.1, pp. 69–111.

Tomer, Adie and Joseph Kane (2014). “Mapping freight: the highly concentrated nature of goods trade in the United States”. *Global Cities Initiative Working Paper*.

Van Marrewijk, Charles et al. (1997). “Producer services, comparative advantage, and international trade patterns”. *Journal of International Economics* 42.1-2, pp. 195–220.

Vartia, Yrjö O (1976). “Ideal log-change index numbers”. *Scandinavian Journal of statistics*, pp. 121–126.

Verboven, Frank (1996). “The nested logit model and representative consumer theory”. *Economics Letters* 50.1, pp. 57–63.

Vergriete, Patrice (2013). “La ville fiscalisée: politiques d’aide à l’investissement locatif, nouvelle filière de production du logement et recomposition de l’action publique locale en France (1985-2012)”. PhD thesis.

Weingast, Barry R, Kenneth A Shepsle, and Christopher Johnsen (1981). “The political economy of benefits and costs: A neoclassical approach to distributive politics”. *Journal of political Economy* 89.4, pp. 642–664.

Wilson, Steven G (2012). *Patterns of metropolitan and micropolitan population change: 2000 to 2010*. US Department of Commerce, Economics and Statistics Administration, US …

Wolfers, Justin and Eric Zitzewitz (2004). “Prediction markets”. *Journal of Economic Perspectives* 18.2, pp. 107–126.

Xu, Qiping and Eric Zwick (forthcoming). “Tax Policy and Abnormal Investment Behavior”. *Journal of Finance*.

Zucco Jr, Cesar (2013). “When payouts pay off: Conditional cash transfers and voting behavior in Brazil 2002–10”. *American Journal of Political Science* 57.4, pp. 810–822.

Zwick, Eric and James Mahon (2017). “Tax policy and heterogeneous investment behavior”. *American Economic Review* 107.1, pp. 217–48.