Home Production in the City*

Caterina Soto Vieira†

Job Market Paper
Click HERE for the latest version.

November 5, 2025

Abstract

By 2050, two thirds of the world's population is predicted to live in cities, predominantly mega-cities. Agglomeration drives up productivity, but the commuting costs may act as a brake by incentivizing specialization within the household, and thus potentially decreasing female labor force participation. Using origin-destination travel data from the mega city of São Paulo in Brazil, I document that labor force participation declines sharply with distance from the city center, a spatial gradient largely driven by married women. When women work, they not only face lower wages, but also commute in slower modes of transport, relying more often on public transport and less on driving than men. To quantify the implications of these patterns for the economy, I model the trade-off between benefits of agglomeration and the cost of commuting using a quantitative spatial framework in which couples and singles decide where to live and whether and where to work. My model reveals that if women faced the same commuting costs as men, labor force participation would increase by 10.3 percentage points—an effect larger than equalizing labor market returns by gender—with especially strong impacts for married women on the urban periphery. This shows that investing in transportation infrastructure that makes commuting equal by gender would draw many women into the labor market, substantially narrowing the gender gap in labor supply.

Keywords: Commuting, Gender, Labor Force Participation
JEL Classification: J16, J22, R13, R23, R541, O18

*I gratefully acknowledges the support and guidance of Oriana Bandiera, Gharad Bryan, Robin Burgess, and Daniel Sturm. For very helpful feedback and discussions I thank Shania Bhalotia, Maarten DeRidder, Gabriela Deschamps, Dave Donaldson, Amen Jalal, Isabela Manelici, Alan Manning, Niclas Moneke, Steve Pichske, Veronica Rappoport, Hugo Reichardt, Patrick Schneider, Stephan Thies, Catherine Thomas, and Sarah Winton. I thank participants of the Development, Applied Micro, and Trade and Urban seminars at LSE, the STEG workshop, and the WECW-EGSE conference. Lucas Poli provided excellent research assistance. This research benefited from financial support from STICERD, Hub for Equal Representation, and STEG, which is gratefully acknowledged.

†Department of Economics, London School of Economics: c.soto-vieira@lse.ac.uk.

1 Introduction

By 2050, over two-thirds of the world's population is expected to live in cities (United Nations, 2018). The concentration of people and firms in urban areas has long been recognized as a major source of productivity gains, as agglomeration economies raise output and wages through dense labor markets and externalities (Rosenthal and Strange, 2004; Glaeser and Gottlieb, 2009; Combes et al., 2012; Baum-Snow et al., 2024). Yet urban density also increases congestion and commuting costs, which may limit the extent to which individuals benefit from these productivity advantages. If commuting costs incentivize specialization within the household, density may impair female employment, lowering aggregate labor force participation. This, in turn, can have broader implications for urban labor markets and aggregate welfare. In this paper, I study whether commuting affects within-city gender gaps in labor supply and its aggregate and distributional consequences.

There are two main forces through which commuting may act as a barrier to female employment. The first operates through differences in the effective return to market work: when commuting costs are high, the net benefit of working outside the home falls, making home production relatively more attractive.[1] The second arises from differences in commuting experiences: women often face higher mobility constraints, such as limited access to safe or fast transport and the need to coordinate travel with care responsibilities. If it takes longer for women than for men to move around the city, this also makes home production relatively more attractive to them than market work.

What are the costs to the economy resulting from these gendered commuting frictions? Understanding their aggregate costs requires a framework in which men and women trade off the benefits of agglomeration with the costs of commuting, allowing for specialization when the household is a couple. It should also incorporate how firms and individuals compete for floor space in the city, and how prices, both wages and house prices, respond to these demands.

To formalize these mechanisms, I develop a quantitative spatial model in which men and women, single or in couples, inhabit a city and choose where to live and whether and where to work. They consider the wages, commuting costs, housing prices, and

[1] Throughout the paper, I refer to home production as the alternative to market work. I do not consider home production in the intensive margin; rather, only in the extensive margin, such as in Hsieh et al. (2019) and Caliendo et al. (2019).

amenities associated with their location and labor supply choices. The model allows for endogenous labor force participation and within-couple specialization, meaning that one partner may optimally withdraw from the labor market when commuting costs are high. In this sense, commuting costs act as a force for specialization: because couples take into account both partners' commuting costs, it can be less costly for the household if one member stays at home while the other works.

In the model, individuals who work receive a wage and incur a commuting cost, which depends on both the travel time between their home and workplace and on a commuting elasticity that captures how sensitive they are to increases in travel time. This framework allows men and women to differ in their commuting costs in two ways: first, through differences in their elasticity of commuting flows with respect to travel time; and second, because they may face different effective travel times across the city. In addition, workers benefit from a workplace-specific amenity associated with the neighborhood in which they are employed.

I examine the mechanisms of the model in the context of the mega city of São Paulo, Latin America's largest city. It generates roughly one-tenth of Brazil's GDP but also has among the longest average commuting times by public transport in the world (see Appendix Figure A2).[2,3] By 2010, it was home to over 11 million inhabitants, many of whom are out of the labor market. The city offers an example of the coexistence of high productivity and large spatial inequalities in access to jobs. Moreover, the city's Origin-Destination Surveys (*Pesquisa Origem e Destino*) are ideal for understanding intra-household behavior, as they provide detailed individual, household, and trip information for over 125,000 people, 43,000 households, and 250,000 trips.

Using these surveys, I document several stylized facts on labor supply and commuting by gender and marital status. Labor force participation declines sharply with distance from the city center, where most jobs are concentrated. This pattern likely reflects sorting by income, as individuals outside the labor market tend to live in cheaper neighborhoods, with lower access. Controlling for individual characteristics such as wealth quintile and education largely eliminates this spatial gradient for singles and married

[2] The municipality of São Paulo accounted for 10.3% of Brazil's GDP in 2019. Source: https://agenciadenoticias.ibge.gov.br/en/agencia-news/2184-news-agency/news/32610-sao-paulo-accounted-for-10.3-of-brazilian-gdp-in-2019.

[3] According to the *Moovit Public Transport Index* (2022), São Paulo ranks 10th in the world, with an average travel time by public transport of 62 minutes. Source: https://moovitapp.com/insights/en/MoovitInsightsPublicTransitIndex-cities.

men.[4] However, it remains pronounced for married women, consistent with household specialization.

I then examine gender differences in commuting among those who work. Using reported travel times and distances from the survey, I find that womens commutes take, on average, 5% longer than men's while being 7% shorter in distance, conditional on origin and destination fixed effects. This means that even when men and women live in the same neighborhoods and work in the same destinations, women face longer travel times. When I control for individual characteristics and distinguish between singles and couples, I find no gender differences in commuting distance among singles, but married women are significantly more likely to work closer to home than their husbands. These differences in distance are not mirrored in commuting time, which remains similar across gender and marital status. The pattern is largely explained by transport mode: women, both single and married, are less likely to drive and more reliant on public transport, which is slower and less flexible. This suggests that married women work closer to home not because they value proximity more, but because traveling the same distance is effectively more costly for them. Single women, for whom time constraints are less binding, are willing to commute farther and longer, despite having similar car usage rates.

Labor market returns also differ systematically by gender and marital status. Even after controlling for occupation, sector, and education, large wage gaps remain between men and women. I interpret these residual wage wedges as arising from discrimination in the labor market, which further lowers the effective return to market work for women, particularly among those facing higher commuting costs.

Building on the stylized facts, I structurally estimate the model to quantify how commuting costs and labor market discrimination shape labor supply and spatial sorting within the city. The first step is to estimate gravity equations for couples and singles to recover some of the models parameters, such as the commuting elasticities and the spatial differences in the amenity of home production in each neighborhood, which I call the home sector pull. Identification comes from distinct sources of variation: commuting elasticities are identified from how commuting flows respond to travel times, and home production amenities from differences in participation rates across neighborhoods.

[4] Throughout the text, I use "married" and "in a couple" interchangeably to refer to heterosexual couples living together, whether formally married or not. All other working age individuals are considered single.

Importantly, the gravity equation for couples can be estimated by considering the commuting time for husband and wife simultaneously, under the assumption that commuting costs enter household utility in a multiplicative way. This functional form is key to generating the specialization result within the household, fits the patterns present in the data very well, and grants tractability to estimate the model. The travel time matrices for men and women, single or married, are a weighted average of the travel times by public transport, cars, and walking. The weights reflect the likelihood that men and women will use each mode of transport in the survey.[5] The results of the gravity equation indicate that men are less sensitive to commuting time than women are, with smaller elasticities of commuting flows with respect to time in absolute value.

With the commuting elasticities and the home-production pull parameters at hand, the estimation proceeds by recovering location-workplace choice probabilities from the data using observed residence and workplace locations by gender and marital status. In the model, home production is treated as another workplace. I then invert the model and recover the fundamental parameters, such as productivity and amenities.

With the estimated model, I conduct two counterfactual exercises to assess how gender differences affect the distribution of paid and unpaid economic activity. In the first, I remove labor market discrimination by equalizing men's and women's wages at each location. The increased return to market work induces an aggregate increase in labor supply of 9 p.p., from a baseline in which 23.2% of the city's population are out of the labor market. The effects are mostly driven by married women moving from home production to market work. In contrast, married men increase their rates of home production due to specialization forces within couples. Despite the large changes in labor supply, there is no effect on the spatial gradient of home production with respect to distance from the city center.

In the second counterfactual, I equalize commuting conditions between men and women, allowing them to face the same travel times and commuting sensitivities. This would decrease home production rates in the city by 10.3 percentage points, a slightly larger increase relative to equalizing wages. Unlike the previous counterfactual, removing commuting differences is more effective at decreasing home production in the periphery, with large impacts on the spatial gradient of home production in the city.

While both counterfactuals raise aggregate output through higher labor supply, the

[5] Tsivanidis (2022) and Pietrabissa (2023) compute the travel time matrices for high-skilled and low-skilled using the same approach.

increase is larger when commuting differences are removed. Reducing spatial frictions allows individuals to access the most productive areas of the city, amplifying the productivity gains from greater participation. By contrast, welfare gains are more muted when wages are equalized. Higher nominal incomes boost housing demand, especially in central locations, driving up floor space prices and offsetting part of the real income gains from eliminating discrimination. General-equilibrium adjustments thus absorb a share of the nominal wage increase, as higher earnings are capitalized into housing prices rather than higher consumption. Equalizing commuting costs, in turn, delivers much larger welfare improvements by directly reducing travel disutility and flattening the housing price gradient, enabling households to reach good jobs without paying for central area's housing prices.

Importantly, equalizing commuting conditions is not necessarily equivalent to improving transport infrastructure. Expanding public transport networks or reducing congestion may shorten average travel times, but such improvements do not necessarily close gender gaps in commuting if women remain less likely to drive or continue facing safety risks in public transit. Policies that improve infrastructure without addressing these gender-specific barriers may therefore have limited effects on labor market participation and aggregate output.

Finally, in light of the two counterfactual results, it is important to highlight that reducing differences in commuting between men and women may be easier to achieve in practice than fully eliminating discrimination in the labor market. Yet, in the model, both policies yield comparable gains in participation and aggregate output. This suggests that targeted interventions that reduce gendered commuting frictions can be a powerful tool to promote gender equality in urban labor markets, even in the presence of persistent wage discrimination.

Related literature A large literature documents gender differences in both the willingness to commute and realized commuting patterns. Women tend to accept shorter commutes and value commute reductions more than men, trading off wages for proximity (White, 1986; Le Barbanchon et al., 2021; Petrongolo and Ronchi, 2020; Borghorst et al., 2024; Bütikofer et al., 2025).[6] Much of this literature treats commuting time and distance as interchangeable measures of commuting costs. Using data that jointly report

[6] Oreffice and Sansone (2023) show that gender commuting differences are present mostly in same-sex couples, suggesting that gender-conforming norms play an important role.

commuting time, distance, and main mode of transport, I show that this assumption is misleading. In the data, gender gaps are smaller for time than for distance, a pattern consistent with women being less likely to drive and more likely to use public transport or walking/cycling–modes that typically cover shorter distances for a given travel time.[7]

A second strand of work emphasizes the effects of barriers to women's mobility, often emphasizing safety and monetary costs. I show that differences in access to faster modes of transport can impact outcomes of women. Prior studies highlight the role of monetary costs (Dasgupta and Datta, 2023; Chen et al., 2024) and safety concerns (Borker et al., 2021; Garlick et al., 2025) as key obstacles. These constraints shape women's educational outcomes (Cheema et al., 2019; Borker et al., 2021; Fiala et al., 2022; Alba-Vivar, 2024), employers hiring behavior (Buchmann et al., 2024), and women's job-switching decisions (Sharma, 2023). In Mexico City and São Paulo, women exhibit a positive willingness to pay for women-only subway cars, consistent with safety risks raising commuting costs (Aguilar et al., 2021; Kondylis et al., 2025), and Christensen and Osman (2023) find that discounted Uber rides in Cairo increase female mobility, largely through safety improvements.[8]

A third and closely related literature examines how commuting frictions shape women's labor market outcomes. Much of this work studies cross-city differences in female labor force participation in the US (Black et al., 2014; Moreno-Maldonado, 2022; Farré et al., 2023). Evidence from China, a rapidly urbanizing economy, also shows that rising travel times explain a large share of the recent decline in women's labor force participation (Liu et al., 2025). This paper is closer to research that explores this relationship within cities (Kawabata and Abe, 2018; Liu and Su, 2022; Velásquez, 2022; Ranosova, 2025).[9] The contribution is to embed these mechanisms in a general equilibrium framework with endogenous location and labor supply decisions. Through counterfactual analysis, I study how removing gender differences in commuting affects choices between home production and market work. In practice, this could occur through improved transportation infrastructure, as shown in studies finding positive effects of new transport links on women's labor supply (Seki and Yamada, 2020; Velásquez, 2022; Caldwell and Danieli,

[7] Le Barbanchon et al. (2021) show that the gender gap in reservation wages and maximum acceptable commute varies by geography and suggest that car dependence in some locations may be a key explanation.

[8] For a thorough review of the literature on barriers to female autonomy and labor market participation, see Heath et al. (2024) and Jayachandran (2020).

[9] Joshi (2024) shows that commuting affects women's choice between agricultural and non-agricultural jobs.

2024).[10,11] An important caveat, however, is that infrastructure improvements alone do not guarantee gender-neutral commuting conditions, as emphasized in the counterfactual exercises in this paper.

This paper builds on the classic theories of time allocation and household specialization developed by Becker (1965) and Mincer (1962), in which the returns to market work depend on wages and participation costs. High commuting costs lower the effective return to market work, reinforcing specialization: the partner with higher potential wages commutes to work while the other stays home to produce without incurring those costs. Carta and De Philippis (2018) and Ranosova (2025) also link commuting to household specialization. My paper builds on and extends these papers by considering a general equilibrium setting, capturing how changes in labor supply feed back into agglomeration and local prices.

A large literature quantifies the aggregate output losses from gender inequality. Declining barriers in occupational choice and human-capital accumulation account for a substantial share of U.S. growth (Hsieh et al., 2019), and related work finds sizable productivity gains from narrowing gender gaps in participation and entrepreneurship (Cavalcanti and Tavares, 2016; Cuberes and Teignier, 2016). Ngai and Petrongolo (2017) show that the marketization of home production increases aggregate output by reallocating women's time toward market activities, while Bell et al. (2019) document large innovation losses from unequal access to opportunity. My contribution is to show that gender differences in commuting costs and mobility constraints can have aggregate consequences, by shaping women's labor-force participation in urban economies where most employment and output are concentrated.

My paper also relates to the literature on quantitative spatial models, which provide structural frameworks to study how spatial frictions shape the organization of economic activity across space. Most models are designed to explain the spatial distribution of *paid* economic activity, where firms locate and where workers supply market labor, while abstracting from the geography of *unpaid* economic activity, such as home production. Canonical contributions formalize how productivity, amenities, and commuting costs shape the equilibrium allocation of employment and income across space (Allen and

[10]Infrastructure effects are not always positive for women; Bütikofer et al. (2024) show that gaining access to a larger labor market increased the gender wage gap.

[11]Relatedly, other spatial policies that increase female labor supply include housing (Franklin, 2020), childcare (Hojman and Boo, 2022), and working from home (Ho et al., 2024; Jalota and Ho, 2024).

Arkolakis, 2014; Ahlfeldt et al., 2015; Monte et al., 2018), and subsequent work has systematized these insights into general frameworks for quantifying urban structure and welfare (Redding and Rossi-Hansberg, 2017; Redding, 2023, 2024). I extend the framework by introducing home production as an alternative to market work, which endogenizes labor supply decisions, in a similar fashion to Hsieh et al. (2019) and Caliendo et al. (2019).

Yet, to understand the spatial distribution of home production and non-market work, it is necessary to model decisions made within the household. I focus on couples, for whom location and labor supply choices are jointly determined. A growing set of papers embeds couples into quantitative spatial frameworks to study joint residential and workplace decisions within cities (Velásquez, 2022; Ranosova, 2025), showing that commuting costs and workplace heterogeneity generate meaningful gender and welfare implications.[12] Relative to these papers, I allow labor supply endogenous for all individuals, so that commuting frictions influence whether and where individuals work. This extension links the city's spatial structure directly to the extensive margin of labor supply and to the distribution of home production across neighborhoods. Furthermore, I consider that commuting times of men and women may differ due to access to faster modes of transport, considering another way in which commuting can be gendered.

The role of household structure in location choices has received increasing attention in the quantitative spatial literature. Across cities, Fan and Zou (2021) integrate local labor and marriage markets to show how the joint search for jobs and partners shapes the spatial distribution of skills, while Moreno-Maldonado (2022) considers how city size interacts with female labor supply choices. Within cities, Moreno-Maldonado and Santamaria (2022) considers how delayed fertility influenced the urban revival in US cities, and Ahlfeldt et al. (2025) study how family type and age impact residential sorting within the city of Copenhagen.

The remainder of the paper is structured as follows. Section 2 presents the theoretical framework. Section 3 describes the data and context, situating the models mechanisms within the mega-city of São Paulo. Section 4 describes the estimation of the model. Section presents the counterfactual results, where I quantify the aggregate costs of gendered commuting. Section 6 concludes.

[12] Moreno-Maldonado (2022) develops a quantitative framework in which couples decide whether to reside in small or large cities.

2 A Quantitative Spatial Model with Home Production

I develop a quantitative urban model which accounts for the spatial sorting of families in space and for their labor supply choices. In this model, men and women inhabit the city, and they may live in households as single or as couples.[13] The city is formed by different neighborhoods that will vary in terms of their productivity in the production of a tradable good, and in terms of amenities. Amenities make neighborhoods more or less attractive to live in. Neighborhoods can also vary in terms of their home-production pull, which allows the benefit of staying out of the labor market to vary across space. This would capture for instance different gender norms around the city, or access to amenities that are complementary to home production.

2.1 Model Setup

I consider a city with a discrete set of locations $\Theta = \{1, ..., N\}$, which in my data are the finite set of neighborhoods. There are $N+1$ workplaces $\Theta_E = \{0, 1, ..., N\}$, where workplace indexed by zero indicates home production. The city is populated by men and women indexed by $g \in \{m, f\}$, respectively. Households are indexed by $o = \{s, c\}$ after singles and couples. Therefore, there are four types of individuals, and three types of households. All households choose residence and workplace(s) simultaneously. For couples, there is a trilateral choice for the residence indexed by n, the husband's workplace indexed by i, and the wife's workplace indexed by j. Single households choose residence and one workplace location as in Ahlfeldt et al. (2015). Individuals og may differ in terms of their commuting costs, their travel time matrix, and the returns they obtain from market work.

I consider a closed city model, with a fixed number of individuals in each group L^{og}. Locations differ in terms of productivity in the good sector A_n, and amenities B^{og}_{nij}. Firms use labor and floor space to produce a market good m_{nij}, which is the numeraire in the economy with $P = 1$. Individuals out of the labor market are equally productive in producing home goods, which are perfectly substitutable with market goods, and are consumed within the household. They also receive a transfer from the government (e.g. unemployment insurance or cash transfers). All markets are competitive. For conciseness, the model is derived only for couples here, see Appendix Section **??** for the complete version with couples and singles.

[13] The model only considers different-sex couples, and there is no endogenous formation of couples.

2.2 Preferences

Consider a household c with residence in location n, with the husband's workplace in location i and the wife's workplace in location j. Let $\phi_{ij} \in \{0, 1, 2\}$ indicate how many household members are in home production. The utility of household ω depends on consumption of goods c_{nij} and the consumption of residential floor space h_{nij}, location amenities B^c_{nij}, and commuting costs from residence n to workplaces i and j, d^{cm}_{ni} and d^{cf}_{nj}.[14] Preferences are Cobb-Doublas and couples spend a share $1 - \beta^c$ of their income on floor space:

$$U^c_{nij}(\omega) = \frac{B^c_{n0} B^c_{nij} z^c_{nij}(\omega)}{d^{cm}_{ni} d^{cf}_{nj}} \left(\frac{c_{nij}}{\beta^c}\right)^{\beta^c} \left(\frac{h_{nij}}{1 - \beta^c}\right)^{1-\beta^c}, \quad \beta^c < 1, \tag{1}$$

where the iceberg commuting cost $d^{cg}_{ns} = e^{\kappa^{cg} t^{cg}_{ns}} \in [1, \infty)$ increases with the travel time individual g takes to commute between locations n and s $\left(t^{cg}_{ns}\right)$, allowing also different decays by individual. There is a household specific idiosyncratic shock $z^c_{nij}(\omega)$, which captures the household's location decision in terms of residence and workplaces, considering the home sector as an additional workplace location. Individuals may have different productivities the labor market, resulting in different wages by individuals when working for firms, and zero wages otherwise ($w^{og}_0 = 0$). The household maximizes utility subject to a budget constraint:

$$m_{nij} + h_{nij} q_{Rn} \leq (1 - \tau)(w^{cm}_i + w^{cf}_j) + \phi_{ij} T \tag{2}$$

where the price of residential floor space is given by q_{Rn} and the price of the market good is $P = 1$. Individuals that work for firms are taxed by the government with a tax rate τ to fund an exogenous transfer T to individuals out of the labor market. This transfer allows households to pay for market goods and housing, even if they don't have wage income. Perfect substitutability between home and market good implies that the

[14] The location amenity for couples is composed of a residential term, a workplace term for each household member, and a home-production pull for each residence location. The full expression is $B^c_{nij} = B^c_{Rn} B^{cm}_{Fi} B^{cf}_{Fj} B^{cm}_{n0} B^{cf}_{n0}$.

shadow price of home goods equals that of market goods.

$$m_{nij} + h_{nij}q_{Rn} \leq (1-\tau)(w_i^{cm} + w_j^{cf}) + \phi_{ij}T \tag{3}$$

$$c_{nij} + h_{nij}q_{Rn} \leq (1-\tau)(w_i^{cm} + w_j^{cf}) + \varphi\mathbb{1}(j=0) + \varphi\mathbb{1}(i=0) + \phi_{ij}T \tag{4}$$

Implicitly, home produced goods increase the effective budget constraint by allowing the household to increase expenditure by how much they save on goods when producing them at home. The optimisation of utility leads to the optimal demands of the goods and housing, conditional on the residential location and the location of workplaces:

$$c_{nij} = \beta^c((1-\tau)(w_i^{cm} + w_j^{cf}) + \phi_{ij}(\varphi + T) \tag{5}$$

$$h_{nij} = \frac{1-\beta^c}{q_{Rn}}((1-\tau)(w_i^{cm} + w_j^{cf}) + \phi_{ij}(\varphi + T)) \tag{6}$$

Let the household total income be defined as $HI_{ij}^c = (1-\tau)(w_i^{cm} + w_j^{cf}) + \phi_{ij}(\varphi + T)$. Substituting the demands on the utility function, I obtain the indirect utility function of couples:

$$V_{nij}^c(\omega) = \frac{B_{nij}^c HI_{ij}^c z_{nij}^c(\omega)}{d_{ni}^{cm} d_{nj}^{cf} q_{Rn}^{1-\beta^c}} \tag{7}$$

which depends on the household total income, the price of floor space, location amenities, labor market participation amenities, commuting costs of household members, and the household idiosyncratic shock. The trade-off between doing working or not is captured in the indirect utility by the different returns and amenities couples receive for each choice. Market work increases household income, but it incurs commuting disutility, unlike the home sector, which has a fixed return and no commuting costs ($d_{n0}^{cg} = 1$). An important feature that allows for specialization in couples, with one person in market work with the other in home sector, is the multiplicative form of the commuting costs of husband and wife. In practice, it augments the agglomeration force for working couples, and it make specialization attractive in the peripheries of cities, where commutes are longer. Finally, I assume that for each couple household, the idiosyncratic component of utility z^c_{nij} is drawn from an independent Fréchet distribution:

$$F(z) = e^{-z^{-\epsilon^c}}, \quad \epsilon^c > 1, \tag{8}$$

This implies that the household's location and labor supply decisions are made jointly and simultaneously, and are governed by the same elasticity. It seems natural that individuals choose where to live and whether to work or not simultaneously, and that these decisions are not independent.[15]

2.3 Residence and Workplaces Decision

Couples choose the residence and workplaces to maximize their utility. Since z^c_{nij} are Fréchet-distributed, the share of households in each residence and workplace(s) presents a closed form solution:

$$\pi^c_{nij} = \frac{\left(B^c_{nij} HI^c_{ij}\right)^{\epsilon^c} \left(d^{cm}_{ni} d^{cf}_{nj} q^{1-\beta^c}_{Rn}\right)^{-\epsilon^c}}{\sum_{r=1}^N \sum_{s=0}^N \sum_{p=0}^N \left(B^c_{rsp} HI^c_{sp}\right)^{\epsilon^c} \left(d^{cm}_{rs} d^{cf}_{rp} q^{1-\beta^c}_{r}\right)^{-\epsilon^c}} \equiv \frac{L^c_{nij}}{L^c} \tag{9}$$

where L^c_{nij} is the measure of couples c living in n and with workplaces i for the husband and j. Therefore, couples sort across residence and workplaces depending characteristics of these location, and their own idiosyncratic preferences.[16] Summing across all possible workplace pairs, I obtain the probability that a couple may choose to live in location n:

$$\pi^c_{Rn} = \frac{\sum_{s=0}^N \sum_{p=0}^N \left(B^c_{nsp} HI^c_{sp}\right)^{\epsilon^c} \left(d^{cm}_{ns} d^{cf}_{np} q^{1-\beta^c}_{Rn}\right)^{-\epsilon^c}}{\sum_{r=1}^N \sum_{s=0}^N \sum_{p=0}^N \left(B^c_{rsp} HI^c_{sp}\right)^{\epsilon^c} \left(d^{cm}_{rs} d^{cf}_{rp} q^{1-\beta^c}_{r}\right)^{-\epsilon^c}} \equiv \frac{L^c_{Rn}}{L^c} \tag{10}$$

If I sum across all the possible residence locations, I find the probability that a couple chooses the pair of workplaces ij:

$$\pi^c_{Fij} = \frac{\sum_{r=1}^N \left(B^c_{rij} HI^c_{ij}\right)^{\epsilon^c} \left(d^{cm}_{ri} d^{cf}_{rj} q^{1-\beta^c}_{r}\right)^{-\epsilon^c}}{\sum_{r=1}^N \sum_{s=0}^N \sum_{p=0}^N \left(B^c_{rsp} HI^c_{sp}\right)^{\epsilon^c} \left(d^{cm}_{rs} d^{cf}_{rp} q^{1-\beta^c}_{r}\right)^{-\epsilon^c}} \equiv \frac{L^c_{Fij}}{L^c} \tag{11}$$

[15]This is equivalent to having a nested Fréchet structure, where the three location choice elasticities are the same.

[16]The focus of the model is on the location decision of couples, and I do not add any other worker heterogeneity to the model. Sorting by skill and/or income are also relevant in cities (Tsivanidis, 2022).

Finally, the probability of choosing workplaces i and j conditional on residence location n is:

$$\pi^c_{nij|n} = \frac{\left(B^c_{n0} B^c_F HI^c_{ij} / d^{cm}_{ni} d^{cf}_{nj}\right)^{\epsilon^c}}{\sum_{s=0}^N \sum_{p=0}^N \left(B^c_{n0} B^c_F HI^c_{sp} / d^{cm}_{ns} d^{cf}_{np}\right)^{\epsilon^c}} \equiv \frac{L^c_{nij}}{L^c_{Rn}} \tag{12}$$

and it depends on workplace characteristics, the commuting costs between residence and workplaces and the participation amenity at the residence location.

2.4 Production

Production of Market Goods. All firms operate in perfectly competitive markets and exhibit constant returns to scale. Output price is normalized to one. Firms combine labor and commercial floor space with a Cobb–Douglas technology:

$$Y_\iota = A_\iota \, L^{\alpha}_{F\iota} \, H^{1-\alpha}_{F\iota}, \qquad \alpha \in (0, 1), \tag{13}$$

where A_ι denotes total factor productivity and α is the labor share. Let $q_{F\iota}$ denote the price of commercial floor space.

Employer Utility Maximization. Following Becker (2010), employers may experience taste-based disutility from hiring certain worker groups cg. Let $\partial^{cg} \geq 0$ denote the utility cost per unit of group cg labor. The employer maximizes utility, rather than profits:

$$\max_{L^{cg}_{F\iota}, H_{F\iota}} U^{owner}_\iota = A_\iota L^{\alpha}_{F\iota} H^{1-\alpha}_{F\iota} - q_{F\iota} H_{F\iota} - \sum_{cg} \left(w^{cg}_\iota + \partial^{cg}\right) L^{cg}_{F\iota}, \tag{14}$$

where $L_{F\iota} = \sum_{cg} L^{cg}_{F\iota}$ denotes total employment in location ι. The first-order conditions imply for all employed groups cg:

$$\alpha A_\iota L^{\alpha-1}_{F\iota} H^{1-\alpha}_{F\iota} = w^{cg}_\iota + \partial^{cg}, \qquad (1-\alpha) A_\iota L^{\alpha}_{F\iota} H^{-\alpha}_{F\iota} = q_{F\iota}, \tag{15}$$

The effective marginal cost of labor is therefore equivalent to the observed wages plus the discrimination cost for the employer.

$$\tilde{w}_\iota = w^{cg}_\iota + \partial^{cg}, \tag{16}$$

Therefore, observed wages incorporate the discrimination wedge

$$w_\iota^{cg} = \tilde{w}_\iota - \partial^{cg} = \left(1 - \gamma^{cg}\right) \tilde{w}_\iota, \qquad \gamma^{cg} = \frac{\partial^{cg}}{\tilde{w}_\iota}. \tag{17}$$

From (15), cost minimization implies the usual unit-cost condition:

$$\tilde{w}_\iota = A_\iota^{1/\alpha}\, \alpha \left(1 - \alpha\right)^{\frac{1-\alpha}{\alpha}} q_{F\iota}^{\frac{\alpha-1}{\alpha}}, \tag{18}$$

and the corresponding demand for commercial floor space:

$$H_{F\iota} = \left(\frac{(1-\alpha) A_\iota}{q_{F\iota}}\right)^{1/\alpha} L_{F\iota}. \tag{19}$$

Goods market clearing. All firms produce the tradable numeraire good c, whose price is normalized to one. Residents of the city are workers only; firm owners live elsewhere. Local households therefore finance their consumption of c with labor income. Firms sell part of their output within the city to these residents and export the remainder. Exports adjust so that the revenue accruing to firm owners is exactly sufficient to drive owner utility to zero in equilibrium. When discrimination is present, lower wages reduce participation and local demand, and a larger share of output is exported.

Home Produced Goods Individuals in home production produce goods that are consumed within the household, and they all have the same productivity in producing these goods. For any individual who works in $\iota = 0$, they produce:

$$\varphi = \begin{cases} \varphi, \text{ if } \iota = 0 \\ 0, \text{ otherwise} \end{cases} \tag{20}$$

Unlike firms, there is no discrimination in home production.

2.5 Commuter Market Clearing

Commuter market clearing requires that the measure of couples that choose workplace pair ij is equal to sum of the measure of couples that live each residence, weighted by the probability of choosing ij conditional on living in n:

$$L^c_{Fij} = \sum_{n=1}^{N} \pi^c_{nij|n} L^c_{Rn} \tag{21}$$

This condition also implies that the expected household income for a couple residing in location n equals the sum of wages in both workplaces, across all workplaces, weighted by the conditional probability of commuting:

$$v^c_n = \sum_{s=0}^{N} \sum_{p=0}^{N} \pi^c_{nsp|n} HI^c_{sp} \tag{22}$$

2.6 Floor Space Market Clearing

I assume that the prices of commercial and residential floor space are such that the markets clear and supply meets the demand. Under this assumption, the price of commercial floor space is given by the first order condition of the firm profit maximisation problem:

$$q_{Fi} = \frac{(1-\alpha)}{\alpha} \frac{\tilde{w}_i L_{Fi}}{H_{Fi}} \tag{23}$$

and the price of residential floor space is given the demands for residential floor space from couples within a location:

$$q_{Ri} = \frac{(1-\beta^c)}{H_{Rn}} v^c_n L^c_n \tag{24}$$

where the the demand for residential floor space in a location depends on the expected income of the households, and the measure of households that choose to reside there.[17]

2.7 Government Budget Constraint

The government collects taxes from paid workers to fund cash transfers to individuals out of the paid labor market. Every working individual pays a flat tax rate of their wage to the government, which then distributes it among all the individuals that are out of

[17] The residential demand is augmented by the singles' demand for floor space when I consider them in the model. This is one of the few changes of adding them, as they compete with the couples for residential floor space.

the labor market. The government's budget constraint is given by:

$$\tau \sum_{g=f,m} \sum_{i=1}^{N} w_i^{cg} L_{Fi}^{cg} = (L_{F0}^{cm} + L_{F0}^{cf})T \tag{25}$$

2.8 Welfare

Welfare is defined as the sum over the utilities of households over all residence locations:

$$V^c = \left(\sum_{n=1}^{N} \sum_{i=0}^{N} \sum_{j=0}^{N} \left(B_{nij}^c H I_{ij}^c \right)^{\epsilon^c} \left(d_{ni}^{cm} d_{nj}^{cf} q_{Rn}^{1-\beta^c} \right)^{-\epsilon^c} \right)^{1/\epsilon^c} \tag{26}$$

2.9 General Equilibrium

Given the model parameters $\{\alpha, \beta^c, \epsilon^c, \kappa^{cg}, \gamma^{cg}\}$, city population L^c, government cash transfers T, and location characteristics $\{A_i, B_{n0}^c, B_{Ri}^c, B_{Fi}^{cg}, t_{ni}^{cg}\}$, the general equilibrium is represented by the vector of endogenous objects $\{L_{Ri}^c, L_{Fij}^c, \tilde{w}_i, q_{Fi}, q_{Ri}, \tau\}$ determined by the following equations:

1. Residential population equation (10).
2. Labor supply equation (11).
3. Zero profit and profit maximisation from firms (**??**).
4. Floor space demand from firms (23).
5. Floor space demand from residences (24).
6. Government budget constraint (25).

2.10 Mechanisms that Lead to Sorting

In the model, three main mechanisms drive the spatial sorting of individuals out of the labor market and toward the periphery of the city. The first and most direct mechanism is that, all else equal, lower-income households tend to locate in cheaper neighborhoods. The second mechanism arises from the fact that individuals who do not participate in the labor market avoid commuting costs and receive a fixed home-production return. For couples living in peripheral areas, where both partners would otherwise face long commutes, this creates strong incentives for specialization, since the household's utility depends on the commuting costs of both members. Finally, spatial sorting may also be driven by local amenities: neighborhoods in the suburbs may offer higher residential amenities that attract households even when the returns to market work are lower due to longer commutes.

The mechanisms behind the gendered specialization of married women out of market work, while their husbands commute to jobs elsewhere, stem from two key asymmetries. First, men and women face different potential returns in the labor market. Second, they experience different commuting costs, shaped by both travel times across the city and gender-specific sensitivities to commuting disutility (κ^{cg}).[18]

3 Data and Context

In this section, I first describe the data sources used for the empirical evidence and for the structural estimation. I follow by describing the context of the city. Finally, I show empirical support for the forces in the model by documenting facts about employment and labor supply by gender and marital status in the city.

3.1 Data

In my empirical and quantitative analysis, I combine data from different sources including household surveys, administrative records, and the national census. I next describe the two main data sources and briefly describe where I obtain additional variables.

Household Travel Surveys The main data source are the household travel surveys for the São Paulo Metropolitan Region for the years 2007, and 2017. Each cross section contains a random sample of households, representative at a small spatial level. The surveys contain information at the household level, describing the household composition, and individual characteristics of all the members, such as education level and activity status. The surveys report the coordinates of residence, workplace, and educational locations, when existing. It also indicates the type of school, sector of work, and occupation.

The survey contains provide detailed information on the daily trips of 126,046 individuals, living in 43,419 households. They report all the trips taken by each household member in the previous day, which is necessarily a weekday. For each trip, the data reports the mode of transport, motive of the trip at origin and destination, reported travel

[18]Gender differences in preferences for home production could also contribute to the observed specialization patterns.

time, with time of departure and time of arrival. Therefore, for each member of the household, I know their activities and measures of mobility.

Travel time matrix To obtain the travel times by mode of transport for each I create a travel time matrix by different modes of transport using a realistic routing algorithm (r5r) for the years 2007 and 2017 (Pereira et al., 2021). Data from Open Street Maps provide the surface roads' speed and direction. These are used by the algorithm to calculate a distribution of travel time by car and walking between any two points in the city. For travel times by public transit, in addition to road surface data, the algorithm requires General Transit Feed Specifications (GTFS) datasets, which provide the schedule, routes, and stops of the public transportation network of the city. This way, it is possible to compute the travel time by public transport between locations.

Additional data I use the National Budget Survey (*Pesquisa do Orçamento Familiar*) to compute the housing expenditure shares, and the Property Tax Registry (ITBI) to obtain floor space prices for both residential and commercial floor space. See Appendix Section A for additional details on the datasets.

3.2 Context

The city of São Paulo is the largest city in Latin America in terms of population, with more than 11 million inhabitants, and serves as Brazil's primary economic center. The municipality alone accounts for roughly one tenth of the country's GDP and concentrates a large share of formal employment, offering a diverse set of job opportunities. Despite this economic importance, it is a city characterized by severe spatial frictions, resulting from a rapid and largely unplanned urbanization process. The average reported commuting time in 2007 was close to 54 minutes, and commutes that rely on public transportation as the main mode exhibit substantially higher average . Many regions of the city remain underserved in terms of infrastructure and services, both privately and publicly provided, which limits accessibility to jobs.

Figure 1 maps the employment and population density quintiles, using the same scale. They show that while just a few central neighborhoods concentrate most employment in the city, the same does not occur for residential choices. Population is much more spread out across different areas of the city, with many peripheral neighborhoods exhibiting large concentration of residents.[19] This is consistent with other cities in the

[19] The extreme south and north of the city have presence of environmentally protected areas, and for

Figure 1: Employment and Population Density in São Paulo

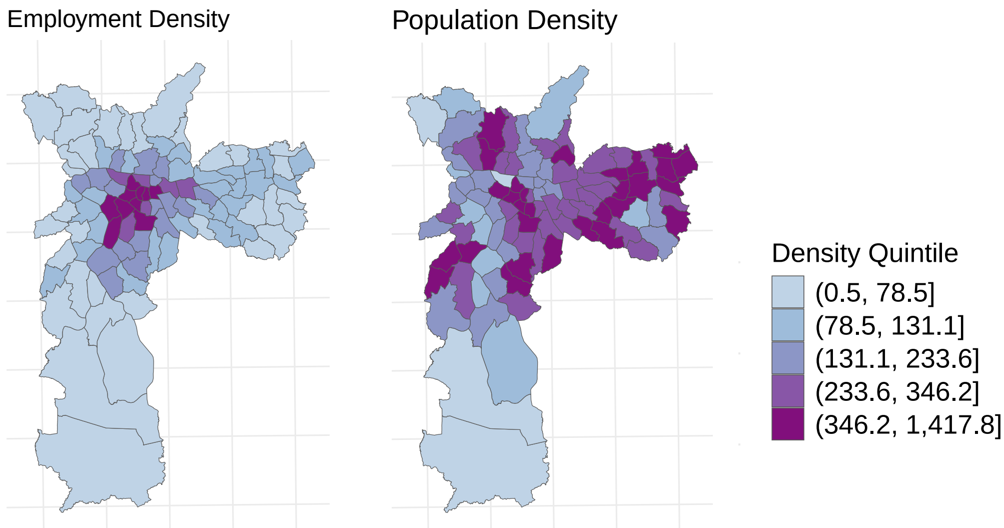

Notes: The map on the left shows the employment density by neighborhood, by counting the total paid employment population and dividing it by the area of each neighborhood, measured in hectares. The map on the right shows the residential population density, as the number of residents divided by the area. There are 96 neighborhoods in total.

world, in which most employment and jobs agglomerate in just a few locations. The center business district is chosen as the Se station, a historical landmark in the city center, and its location is highlighted in Appendix Figure A1.

Furthermore, a large share of the working age population of the city is out of the labour market and gender gaps in participation are pronounced. Appendix Table A1 shows an almost 20 percentage point difference between men and women who report working. This gap is almost entirely explained by differences in women reporting being house wives. The other activities, not working, retired and student, report close to no gender differences. Throughout the paper, I pool together those who report not working with house wives into the home production category. Throughout the remainder of the paper, the sample does not include retirees or students, keeping only those that are working or in home production.

3.3 Stylized Facts

To motivate the quantitative analysis, this subsection documents a set of stylized facts that describe the spatial organization of market and non-market activity within the city of São Paulo. Together, these patterns highlight the spatial frictions and gender asymetries that the model seeks to explain.

that reason have much lower population density.

Figure 2: Share of Individuals in Home Production by Residence Location

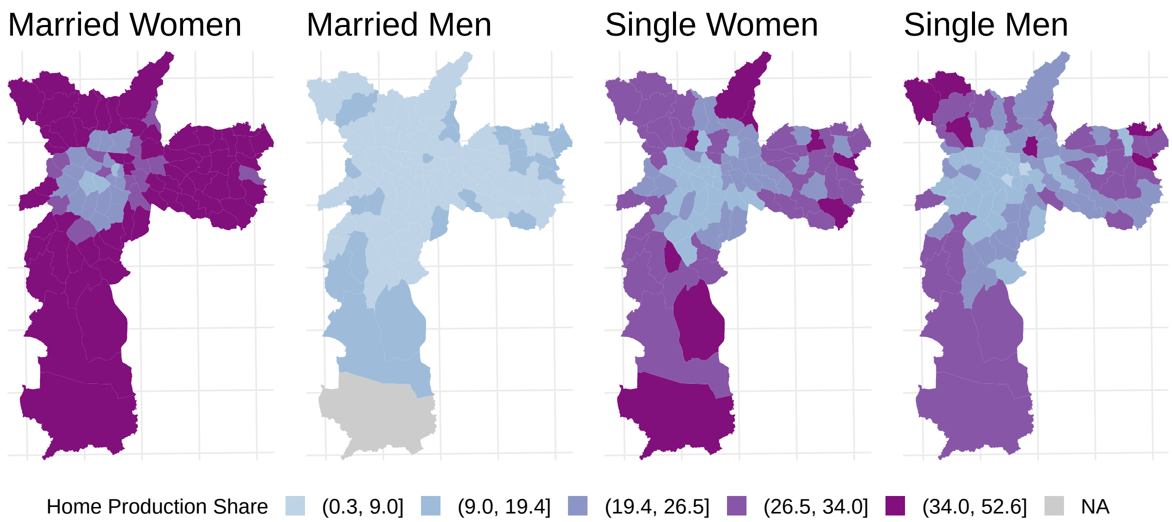

Notes: The figure maps the share of share of working-age individuals engaged in home production across neighborhoods of São Paulo, separately for married women, married men, single women, and single men. Darker shares indicate higher shares of home production. Source: Origin-destination survey, 2007.

Figure 2 maps the share of working-age individuals not engaged in market work across neighborhoods, by gender and marital status. The prevalence of home production varies significantly across space, ranging from less than 5% in the central areas to over 50% in peripheral neighborhoods. However, this distribution is mediated by gender and marital status. Three main patterns emerge. First, the spatial distribution of home production for singles, men or women, is relatively similar. Second, among couples, we see a clear pattern of specialization, with married men rarely doing home production, and their wives exhibiting home production shares of over 34% across most of the city. Finally, for all groups, labor supply is markedly higher in the city center, where job opportunities are.

The patterns observed in Figure 2 could in part be driven by observables. Income and social inequalities are present in the city and they also influence residential sorting (Harari, 2024). The rent gradient would also influence the sorting of individuals with lower resources to peripheral neighborhoods. I address this by regressing the probability of doing home production on distance to the CBD for single and couples. Column (1) of Table 1 shows that a 10km increase in distance from the center is associated with a 5 percentage point increase in the probability of being out of the labor market for single men, and an 8 percentage point increase for single women. Column (2) shows that the gradient is considerably smaller for married men relative to all other individuals. Married women face a seven percentage point increase, similar to single women. In the distance to CBD distribution, 10 km is approximately the difference between being in the 25th to the 75th percentile.

Columns (3) and (4) of Table 1 show the spatial gradients once I control for observable characteristics, such as age, household composition, education level and a measure of wealth quintiles.[20] The results show that adding controls eliminate the distance gradient for single men and halve it for single women, with a coefficient that is only statistically significant at the 10% level. This suggests that for couples the increased likelihood that they will live farther away from jobs if they are out of the labor market can be mostly explained by their characteristics. The same pattern is not present for couples. The gradient differences between married men and women remain the same and statistically significant. The gradient for married men flips sign, suggesting that couples tend to specialize and the husbands' are more likely to work *where* their wives are in home production and vice-versa. This occurs even for the richest, and is also not explained by the presence of children, as the pattern remains for couples without children in the household.

The evidence above points to clear spatial and gender differences in who participates in the labor market. I next look at whether there are differences in outcomes when individuals participate in the labor market. First, I look at how commuting behavior varies across groups, and I run regressions of commuting outcomes on gender and marital status indicators, controlling for individual characteristics such as age, education, household composition, and neighborhood fixed effects. The results of the regression are in Appendix Table A3 and the outcomes of interest are commuting distance, time and whether the main mode of transport is private. Using the effect sizes from the regressions, I calculate the adjusted means relative to the baseline average of the baseline group (single men).

Figure 3a shows the differences in adjusted means, and shows that while married women work 7.9 km away from home on average, married men work 10.8 km away from home, 39% more distant. In contrast, the average commuting distance of single men and women are statistically equal, with 9.4 km and 9 km respectively. These stark

[20] I compute the wealth quintiles based on reported assets in the household, not on earnings. This approach has been used in the DHS surveys, and applied to other survey and census datasets (Smits and Steendijk, 2015; Zipfel et al., 2022; Bandiera et al., 2022).

Table 1: Probability of Home Production by Distance to the CBD

	P(Home Production)			
	Singles (1)	Couples (2)	Singles (3)	Couples (4)
Distance to CBD (10km)	0.05***	0.02***	0.01	-0.01*
	(0.01)	(0.007)	(0.01)	(0.007)
Distance to CBD (10km) × Female	0.03**	0.05***	0.03*	0.05***
	(0.02)	(0.01)	(0.02)	(0.01)
Observations	16,229	16,541	16,229	16,541
Individual-level controls			✓	✓
Household-level controls			✓	✓
Education FEs			✓	✓
Wealth quintile FEs			✓	✓

Notes: Robust standard errors in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Source: SP origin destination survey for 2007. Sample includes adults aged 18 to 55 living in the city of São Paulo who are working or in home production. Controls include age, age squared, household size,number of 0-3 year-olds, 4-10 year olds and number of 75+ in the household in the household. Fixed effects for education level and for wealth quintile based on household assets are included.

differences in commuting distance are not reflected in commuting times. Figure 3b shows that commuting times are long for all groups, with a one-way commute lasting close to 60 minutes. Only single women present an average time statistically different from single men, with commutes lasting 64 minutes. One potential reason why women are commuting the same or longer than men, despite working relatively closer to home, is about *how* they commute.

Figure 3c plots the share who commute to work using a private mode of transport. Women are much less likely to drive than men, regardless of marital status. Across many cities of the world using public transit is slower than driving a car. The share of the area where travel time favors public transportation over car use is estimated to be less than 1% in São Paulo (Liao et al., 2020). Access to a private mode of transport has a clear impact on how fast individuals can go around the city and, consequently, the distances that they are able to cover. Both married and single women are much less likely to drive, but single women still travel similar distances to single men by accepting longer commuting times, unlike married women. This pattern alone, however, does not allow me to distinguish between two possible mechanisms: whether men travel farther because they drive and can access more distant jobs, or whether women travel shorter distances because they rely on slower transport modes and therefore face tighter spatial constraints on where they can work. In the model, the commuting cost depends on time between locations, but also how sensitive individuals are to increases in commuting time. This allows for married and single women to face similar travel time between locations, but have different willingness to commute longer.

Figure 3: Differences in Commuting Outcomes by Gender and Marital Status

(a) Commuting Distance

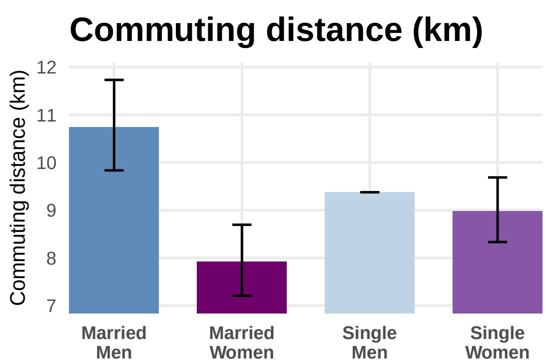

(b) Commuting Time

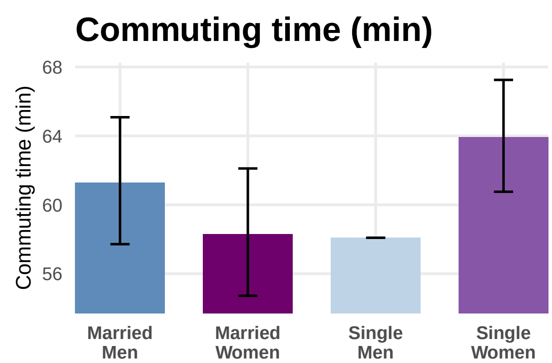

(c) Private Mode Share

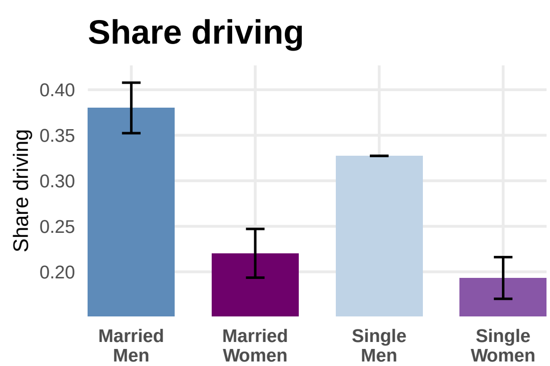

Notes: The figure plots the predicted commuting outcomes (distance, time, and private mode share) for each demographic group, calculated as the baseline mean for single men plus the estimated regression effects for each group. The adjusted means control for age, education, wealth, household composition, and residence location bin fixed effects. The confidence intervals are also derived from the regression estimates (see Appendix Table A3).

Finally, Figure 4 document differences in labor market returns across groups. This is the first channel traditionally emphasized in the literature to explain within-household specialization (Becker, 1965; Mincer, 1962). The higher-earning spouse focused on market work, and the other would allocate more time to home production. After controlling for hours of work, sector, occupation, and education level, married men still earn more

Figure 4: Different Returns in the Labor Market by Gender and Marital Status

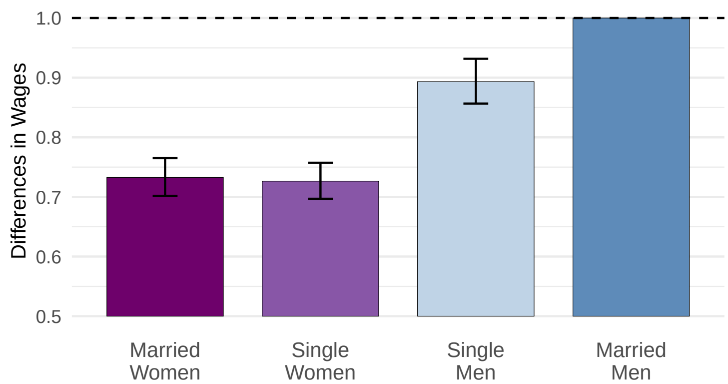

Notes: The figure shows earnings differences across demographic groups, estimated from a log earnings regression controlling for age, educatio, sector, occupation, job type (between regular and flexible), and hours worked on the previous day. The dashed line indicates the reference group (married men) and error bars show 95% confidence intervals. See Appendix Table A4.

than observationally equivalent married women, single women, and single men. While the labor market returns for women are very similar, single men earn less than married men, but still significantly more than women.[21]

While these gaps could also reflect differences in other unobservable characteristics, such as talent, they could also occur due to discriminatory behavior by firms. It is not possible to disentangle these factors, and I choose to interpret them as discrimination, as in Hsieh et al. (2019), as there is no reason to believe that women are systematically less talented than men. The presence of a marital premium for men is consistent with specialization allowing them to devote more time and effort to market work, therefore increasing their market returns.

4 Model Quantification

In this section, I describe how I estimate the quantitative spatial model using data from the city of São Paulo. The main data sources for the estimation are described in Section 3 and the model is estimated at the zone level, and there are 23 zones in the city. The first two estimation steps recover important parameters in the model. I use wage and workplace data to calculate the efficiency units of labor for each individual type in the economy. I run gravity equations to recover the commuting costs semi-elasticities and the participation amenities for the model individuals. The remaining parameters I

[21] This is consistent with evidence in the US showing that there still exist a marriage premium for men, there is no longer a marriage penalty for women Juhn and McCue (2017)

calibrate with other data or I borrow from the literature. With the parameters at hand, I invert the model and back out the economy fundamentals, necessary for the counterfactual exercises.

Discrimination Wedges The discrimination wedges are obtained from a log earnings regression at the individual level, controlling for age, age squared, log hours spent at work the previous day, job type, education level and wealth quintile fixed effects, and sector and occupation fixed effects.[22,23] The regression results are shown in Appendix Table A4 and in Figure 4.

Gravity Equation To estimate the commuting costs elasticities and the home production pull parameters, I rewrite the location choice probability of couples, and apply logs, resulting in the following gravity equation:

$$log(L^c_{nij}) = \underbrace{\eta^{cm}_{n0}}_{B^{cm}_{n0}} + \underbrace{\eta^{cf}_{n0}}_{B^{cf}_{n0}} + \underbrace{\delta_n}_{B^c_{Rn} q^{\beta^c - 1}_{Rn}} + \underbrace{\zeta_{ij}}_{B^{cm}_{Fi} B^{cf}_{Fj} HI^c_{ij}} + \nu^{cm} t^{cm}_{ni} + \nu^{cf} t^{cf}_{nj} + u_{nij} \tag{27}$$

where η^{cm}_{n0} and η^{cf}_{n0} are fixed effects for residence interacted with husband doing home production and with wife doing home production, respectively. I include residence location fixed effects (δ_n), which absorb the residence amenities and floor space prices. Destination pair fixed effects (ζ_{ij}) absorb the workplace amenities for husband and wife and the household income.[24] Travel time between residence and workplaces are given by t^{cm}_{ni} and t^{cf}_{nj}, for the husband and wife, respectively, and they reflect the differences in commuting times that they experience between places due to the different transport modes they use. In Appendix Section A, I describe how I compute the travel time matrices for each individual, and how I validate it with the average reported travel time between locations, available in the origin destination surveys.

The estimated semi-elasticities of commuting are ν^{cm} and ν^{cf}, and they equal $\nu^{cg} = \epsilon^c \kappa^{cg}$. The gravity equation for couples increases the dimensionality and sparsity of the flow matrix, as there are many combination of cells that no couple has chosen. I follow

[22] I classify job types as regular or flexible based on activity status. Regular job accounts for formal employers and employees, civil servants, and sole practitioners, and flexible jobs include informal employers and employees, working in the family business, and self-employed.

[23] The survey does not contain information on how long individuals have been working, and with education level fixed effects, estimated tenure is collinear with age.

[24] The fixed effects also absorb the value of home production ($\varphi + T$), which is the return when in home production.

by estimating the commuting elasticities using Poisson Pseudo Maximum Likelihood (PPML), following Silva and Tenreyro (2006). Furthermore, there can be endogeneity of travel time with respect to flows, if policy makers improve infrastructure where the demand is already high for example (Redding, 2024). In order to address this potential endogeneity, I use distance between the origin and destination as instruments for the travel time using the control-function approach, as suggested by Wooldridge (2015).

The estimated origin-home production fixed effects from the gravity equation correspond to the home production pull parameters in the model, which capture how attractive non-market work is in each residential location. Conditional on workplace characteristics, commuting costs, and residential locations, these parameters are the residual component that rationalizes the observed participation rate of married men and women at each origin. Locations where individuals are more likely to be out of the labor market, even after accounting for these observables, must have a stronger home production pull.

Table 2 shows that the commuting costs for married men are smaller than of married women, indicating that they face lower commuting costs. The absolute value for the semielasticity of married women is 50% larger than of the married men, but they both face relatively small commuting costs, expected due to the long commutes in São Paulo. In Appendix Section X, I show the results of the gravity equation when I consider travel time by car and distance, for both men and women. Using these two other commonly used measures in the gravity equation, they confirm the finding that married women are slightly more sensitive than married men, although with different magnitudes.

The results of the gravity equations for singles are shown in 3. The first takeaway is that, among singles, the commuting cost parameters are very similar across genders, and close in magnitude to those estimated for married women. This suggests that gender differences in commuting elasticities only emerge within couples, when households internalize both partners' commutes and begin to specialize.

Model Inversion With the gravity coefficients and the estimated home production amenities, I proceed by inverting the model. I can rewrite equation (21) from the commuting market clearing condition as a function of only known parameters and employment and residence data to back out the value of $\omega^c = (B^{cm}_{Fi} B^{cf}_{Fj} HI^c_{ij})^{\varepsilon^c}$:

$$L^c_{Fij} = \sum_{n=1}^{N} \frac{B^{cm}_{n0} B^{cf}_{n0} \omega^c_{ij} / exp(\nu^{cm} t^{cm}_{ni} + \nu^{cf} t^{cf}_{nj})}{\sum_{s=0}^{N} \sum_{p=0}^{N} B^{cm}_{n0} B^{cf}_{n0} \omega^c_{sp} / exp(\nu^{cm} t^{cm}_{ns} + \nu^{cf} t^{cf}_{np})} L^c_{Rn} \tag{28}$$

Table 2: Gravity equations of married men and women

D.V.	Total Flow (1)
Commuting Time Husband	-0.024***
	(0.007)
Commuting Time Wife	-0.036***
	(0.007)
Observations	6,788
Residence FE	✓
Workplace Pair FE	✓
Origin-Home Production Wife FE	✓
Origin-Home Production Husband FE	✓

Notes: The table shows the commuting cost semi-elasticities, estimated following the equation (27) with PPML. The unit of observation is residence-workplace of husband-workplace of wife, and the workplaces include home production with travel time equal to zero. Distance between neighborhoods is used as an instrument to travel time in minutes. Own-flows, i.e. commuting within the same neighborhood, are not included in the estimation.

Table 3: Gravity equations of single men and single women

D.V.	Total Flow	
	(1)	(2)
Commuting Time	-0.033***	-0.036***
	(0.005)	(0.005)
Observations	529	529
Residence FE	✓	✓
Workplace FE	✓	✓
Origin-Home Production FE	✓	✓

Notes: The table reports the commuting cost semi-elasticities for single men (column 1) and single women (column 2), estimated using PPML following the same specification as equation (27) but with only one workplace. The unit of observation is residence-workplace, including home production as a potential destination with travel time equal to zero. Distance between neighborhoods is used as an instrument for travel time in minutes. Own-flows (within-neighborhood commuting) are excluded. The origin-home production fixed effects correspond to the estimated home-production pull parameters, B^{sg}_{n0}.

which is a simple system of equations exactly identified. With the vector of ω_{ij}^{c}, I follow the procedure proposed by Ahlfeldt et al. (2015), using an algorithm that exploits the variance of wages in the data to as an additional moment. The algorithm searches for the value of ε^{c} that minimizes the difference between the variance of log wages in the data and the variance of log ω^{c}. Since I don't observe the monetary equivalent return of doing home productiom, I use the minimum wage as the equivalent return. I obtain an estimate of $\varepsilon^{c} = 6.9$ for couples, and lower values for singles, with $\varepsilon^{sf} = 3.57$ for single women and $\varepsilon^{sm} = 3.43$ for single men.

The Fréchet shape parameter governs the dispersion of idiosyncratic preferences in the location-workplace choice. Because couples jointly choose a residence and two workplaces, their choice set is of size $N \times (N+1) \times (N+1)$, whereas singles face $N \times (N+1)$ alternatives. Although the total number of populated cells does not differ drastically—since the couples represent roughly half of the population, and single men and women each a quarter—the dimensionality of the couples' problem implies a much larger potential choice set. Consequently, the model requires a higher Fréchet shape parameter to replicate the observed concentration of commuting flows in a relatively small subset of feasible residence-workplace pair combinations. In other words, couples' decisions appear more deterministic, consistent with the idea that joint household choices exhibit less idiosyncratic dispersion than individual ones.

Data for the average wages for married men at each location are used to construct the effective marginal cost of a unit of labor $\tilde{w}_i$. For individuals subject to firm discrimination, the effective return is adjusted to $(1 - exp(\gamma^{cg}))\tilde{w}_i$, with γ^{cf} being the coefficients from regression shown in Appendix Table A4. Labor market returns are normalized by the geometric mean so they are in the same scale. With known ω, the Fréchet shape parameter, and the monetary returns at each location, it is possible to back out the workplace amenities. Using the share of commercial floor space, the price of commercial floor space, and the effective marginal cost of labor $\tilde{w}_\iota$, I recover productivities from the unit-cost condition in Equation (18).[25] The residence location amenities for each group are backed out from the residence location choice probability stated in Equation (10). Using the commuting probabilites and wages at each location, I compute the expected household income. Lastly, floor space in each location is estimated from the market clearing condition. Further details of the estimation procedure are in the Appendix Section X.

[25] I set the input share of commercial floor space as 20%, consistent with other papers in the literature (Ahlfeldt et al., 2015; Tsivanidis, 2022).

Model Fit and Parameters With the fundamentals of the model, I solve for the baseline equilibrium to see how well the model does at recovering the endogenous variables that are observed in the data. In this solver, I target the aggregate shares of home production by gender and marital status, but not the location-level home production shares, which remain an untargeted moment in the model. I list the parameters in Table 4, and most are estimated together with the model inversion. The return in home production is set to be the value of the minimum wage in 2007, but the results are robust to using the lowest average wage at any location in the city instead. I impose that the transfer represents 30% of this value, but this assumption does not impact the counterfactuals, as I am not considering changes in the tax schedule. This assumption implies that the labor income tax rate is close to 7%.

Table 4: Model parameters, values, and sources

Definition	**Parameter (value)**	**Source**
Household preferences		
Housing expenditure share (couples)	$\beta^c = 0.20$	Calibrated, Budget Survey
Housing expenditure share (singles)	$\beta^s = 0.25$	Calibrated, Budget Survey
Fréchet shape, couples	$\varepsilon^c = 6.9$	Estimated
Fréchet shape, single women	$\varepsilon^{sf} = 3.57$	Estimated
Fréchet shape, single men	$\varepsilon^{sm} = 3.43$	Estimated
Return in Home Production	$(\varphi + T) = R\$415$	Calibrated
Commuting and mobility		
Commuting cost semi-elasticity	$\nu^{cm} = -0.024$	Estimated
Commuting cost semi-elasticity	$\nu^{cf} = -0.036$	Estimated
Commuting cost semi-elasticity	$\nu^{sm} = -0.036$	Estimated
Commuting cost semi-elasticity	$\nu^{sf} = -0.033$	Estimated
Firms and technology		
Labor share in production	$\alpha = 0.8$	Ahlfeldt et al. (2015)
Taste-based discrimination wedge	γ^{cg}	Estimated
Government and transfers		
Transfer to non-participants	$T = 0.3(\varphi + T)$	Calibrated
Labor income tax rate	$\tau = 0.07$	Estimated

Notes: Estimated parameters are obtained from the gravity equations and model inversion. Calibrated parameters are obtained from data, and/or set exogenously. The remaining parameters are obtained from the literature.

Figure 5 compares the predicted and observed values at the location level for employment population, residential population, the share of households in home production, and residential floor space prices, when I solve for the baseline equilibrium. The model provides a good fit to the observed data across the equilibrium outcomes shown below, and others. The tight clustering of points around the 45-degree line indicates that the estimated fundamentals and parameters reproduce the spatial distribution of paid and unpaid economic activity with accuracy. Even for the share of home production in each location, a moment that is untargeted in the estimation, the model is still able to reproduce quite well the distribution of home production in the city by type. Overall, the correlations between the predictions and the data are very high, over 95% across all variables.

Home production pull A key improvement in the model's fit comes from explicitly estimating the home production pull parameters, and including it throughout the estimation. This parameter allows the model to capture spatial heterogeneity in labor market participation behavior, which is not explained by commuting costs or other residence specific amenities, which do not interact with home production. Figure 6a shows the relationship between observed and predicted home production shares at each location for married women (CF), when I estimate the model without the home production pull. The model tends to overestimate home production shares in the periphery, and underestimate them in the center of the city. Figure 6b shows that including this parameter in the estimation improves the model's ability to predict correctly how much home production is done at each location. The correlation increases from 86% to 82%, and the RMSE decreases considerably as well.

This parameter in the data is estimated based on variation in the share of individuals out of the labor market in each location, conditional on commuting times and other residence specific characteristics. Several mechanisms could explain spatial differences in the attractiveness of home production such as childcare accessibility, informal support networks, gender norms, access to other services complementary to home production. Within the model's framework, the locations with highest home production pulls for women and single men are the most central locations. This is because in these locations, even with great access to good paying jobs and low commuting costs, many people still choose to be out of the labor market. Therefore, the model rationalizes that these are locations that make home production relatively more attractive.

Figure 5: Model Fit

(a) Employment Population

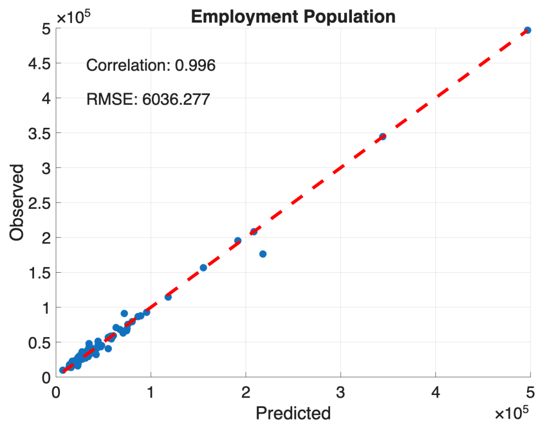

(b) Residential Population

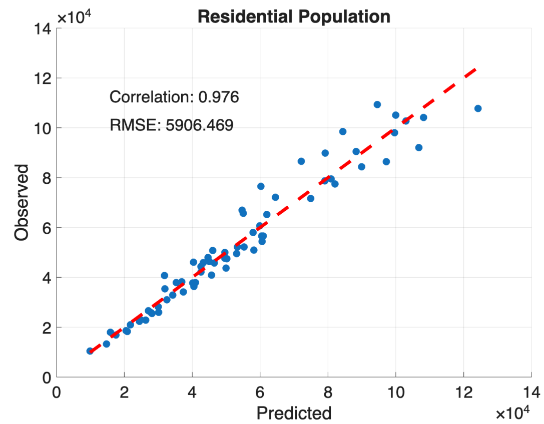

(c) Residential Floor Space Prices

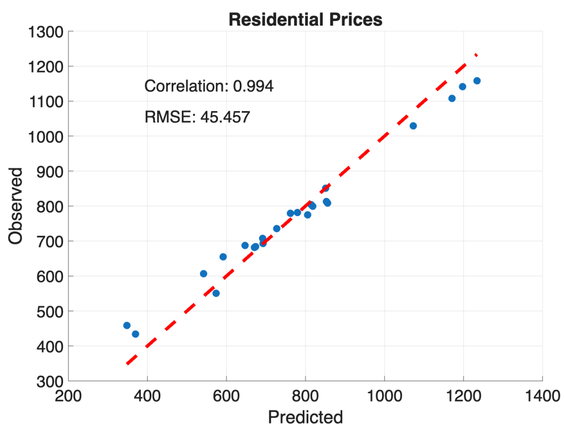

(d) Home Production Share

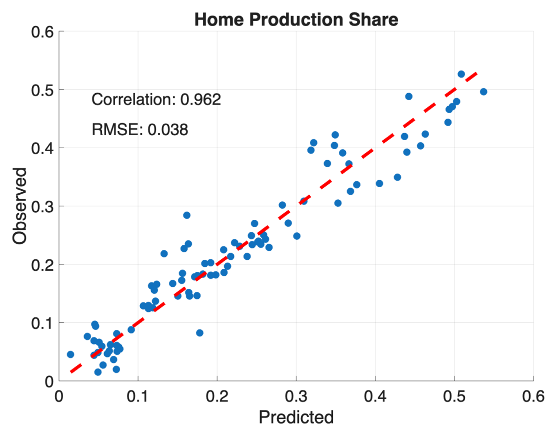

Notes: The figure compares observed and predicted values at the location level for key equilibrium outcomes in the model. Panel (a) plots the employment population, Panel (b) shows the fit for the residential population, Panel (c) displays residential floor space prices and Panel (d) shows the home production share at each location.

Figure 6: Including the Home Production Pull Improves the Model Fit

(a) Share in Home Production without B^c_{n0} (b) Share in Home Production with B^c_{n0}

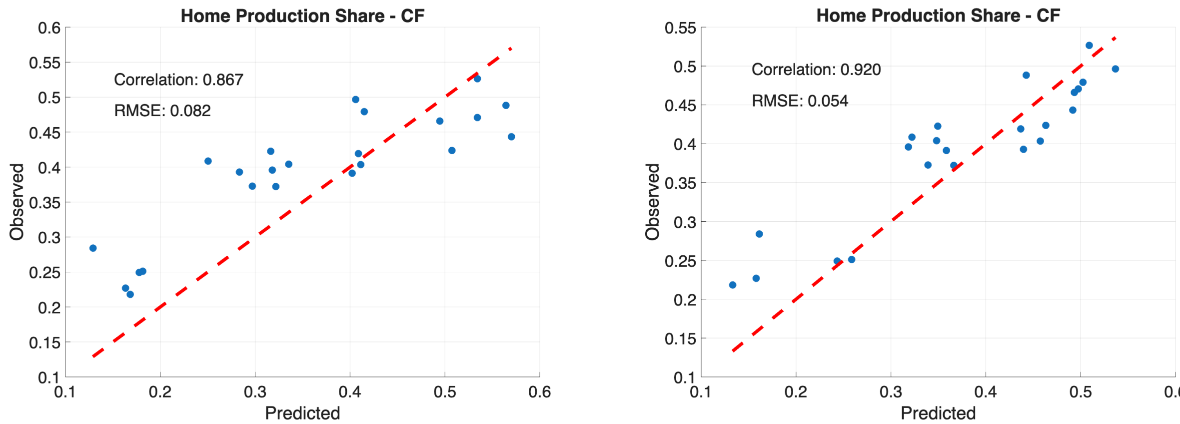

Notes: The figure compares observed and predicted values of home production shares of married men under two alternative estimations. Panel (a) shows the relationship between observed and predicted shares when the model is estimated without home production pull. Panel (b) shows the same graph when the model is estimated considering the home production pull.

5 Counterfactual Exercises

I evaluate two policy-relevant counterfactuals using the estimated model. The first removes gender- and marital-status differences in wages, eliminating taste-based discrimination in the labor market and equalizing effective earnings across individuals working in the same location. The second equalizes commuting conditions across groups, removing gender and marital-status differences in both travel times and commuting elasticities. To do so, I set the travel time matrices and commuting cost semi-elasticities equal to those of a married men for all other individuals. In both exercises, I keep the recovered fundamentals fixed and resolve the spatial equilibrium, allowing wages, labor supply, commuting flows, and floor-space prices to adjust endogenously, while holding housing supply constant in each location.

The model predicts substantial general-equilibrium adjustments under both scenarios, though through different mechanisms. As shown in Figure 7, aggregate output, labor supply, and welfare all increase, but the magnitude and composition of these changes differ across counterfactuals.

When I remove wage differences, firms no longer experience a disutility from hiring women and single men, and the discrimination wedges that previously lowered their wages disappear. Higher wages raise participation, but commuting frictions remain unchanged. Production exhibits diminishing returns to labor and commercial floor space is fixed, therefore an increase in employment translates into a smaller proportional rise in total output. Workers remain constrained in their spatial choices because of commuting frictions.

When I equalize commuting conditions, aggregate output rises by over 12% and moves more closely with employment. Equalizing travel times and commuting elasticities allows individuals to access more productive central jobs, generating both higher participation and a reallocation of labor toward high-productivity workplaces. Even though nominal wages do not change directly, productivity per worker rises because the spatial organization of employment becomes more efficient. As a result, the city achieves larger aggregate output gains than under the wage-equalization counterfactual, with improvements driven by larger labor supply increases, particularly in very productive locations.

Welfare gains differ sharply across the two counterfactuals. When wage differences are removed, higher nominal incomes increase demand for centrally located housing, putting upward pressure on housing prices and partially offsetting real income gains. In contrast, when commuting conditions are equalized, individuals—especially women in couples—can access better jobs at lower travel costs without large changes in nominal wages or housing demand. Because this counterfactual acts directly on commuting disutility rather than through the income channel, reduces inequality in accessibility without generating large rent-capitalization effects. As a result, the city achieves broader welfare improvements.

Removing wage and commuting differences also reshapes household labor supply patterns. In the baseline, the model reproduces the observed specialization of married couples, with men more likely to work in the market and women more likely to engage in home production, especially in peripheral neighborhoods. Figure 8 shows that under the *No Wage Differences* scenario, home production falls sharply for married women. The rise in their effective market returns weakens the incentive for within-household specialization, as the opportunity cost of staying home increases. Married men's home production rises slightly, reflecting re-optimization within couples once both partners' relative wages are equalized. The effects on labor supply for single women are larger than for single men because they suffered more discrimination in the baseline.

Under the *No Commuting Differences* scenario, the decline in married women's home

Figure 7: Counterfactual Changes in Aggregate Outcomes

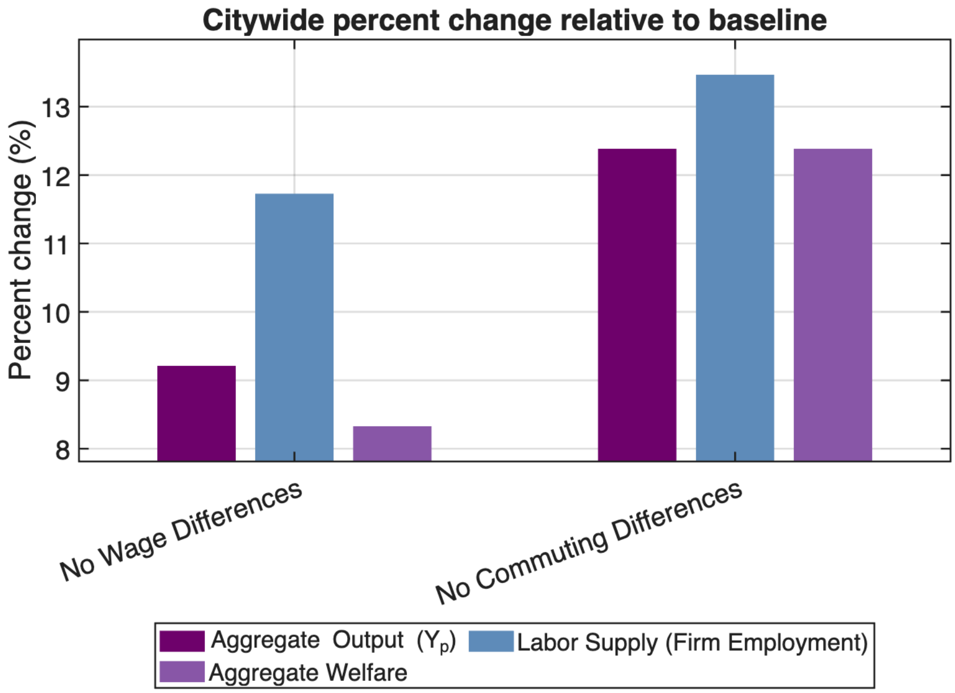

Notes: The bars show the percent change in aggregate output, labor supply, and welfare relative to the baseline level. The first counterfactual equalizes wages between all demographic groups, and the second equalizes commuting costs between all demographic groups.

production is of a similar magnitude to that observed under wage equalization, but it is accompanied by a smaller increase in married men's home production (Figure 8). This pattern arises because, although commuting costs fall for both partners, the returns to market work remain higher for men, limiting the change in their participation decisions. The reduction in commuting frictions nonetheless lowers within-couple specialization overall, as both partners face a more balanced trade-off between engaging in home production and participating in the labor market. For singles, we observe that reducing commuting frictions has a larger effect for single men than eliminating wage wedges, while for single women the effects are stronger under wage equalization due to the role of discrimination (Figure 8). This pattern suggests that single women's participation is more affected by differential returns in the labor market than single men's, although both groups are substantially constrained by commuting costs.

Both counterfactuals also change the spatial gradient of home production across the city. Figure 9 plots the home-production rate of married women against distance to the city center. In the baseline, the share of married women in home production rises

Figure 8: Counterfactual Home Production Shares by Gender and Marital Status

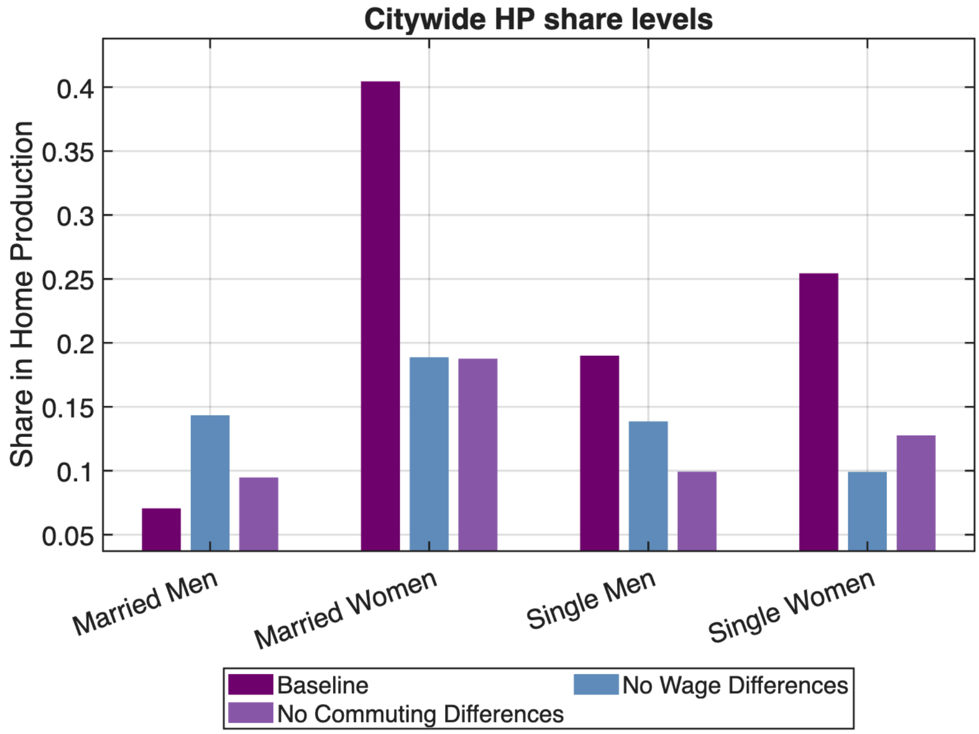

Notes: The figure shows the rates of home production by gender and marital status in baseline and counterfactual scenarios. The dark purple bars show the baseline rates, the blue bars show the counterfactual rates under wage equalization, and the light purple bars show the rates under equal commuting conditions.

steeply with distance to the CBD, reflecting both commuting burdens and spatial sorting of household types. Equalizing wages shifts the level of home production down but does not affect the gradient: peripheral areas remain less attached to the labor market because commuting conditions still differ by gender. Equalizing commuting conditions, by contrast, substantially flattens this gradient. The largest participation gains occur in the periphery, where commuting costs were initially most binding. These results highlight that gendered commuting frictions generate spatially unequal access to market work and reinforce specialization within households.

Figure 10 shows welfare changes by household type. Welfare gains are largest for single women in both counterfactuals, with more modest increases for the other groups. Under the counterfactual without discrimination, single men experience only small welfare gains, as they didn't suffer much discrimination in the baseline. Because wages rise more for women than for single men, the increase in housing prices partially erode the real income of groups that experience smaller nominal gains. When commuting

Figure 9: Counterfactual Home Production Shares by Distance to the CBD

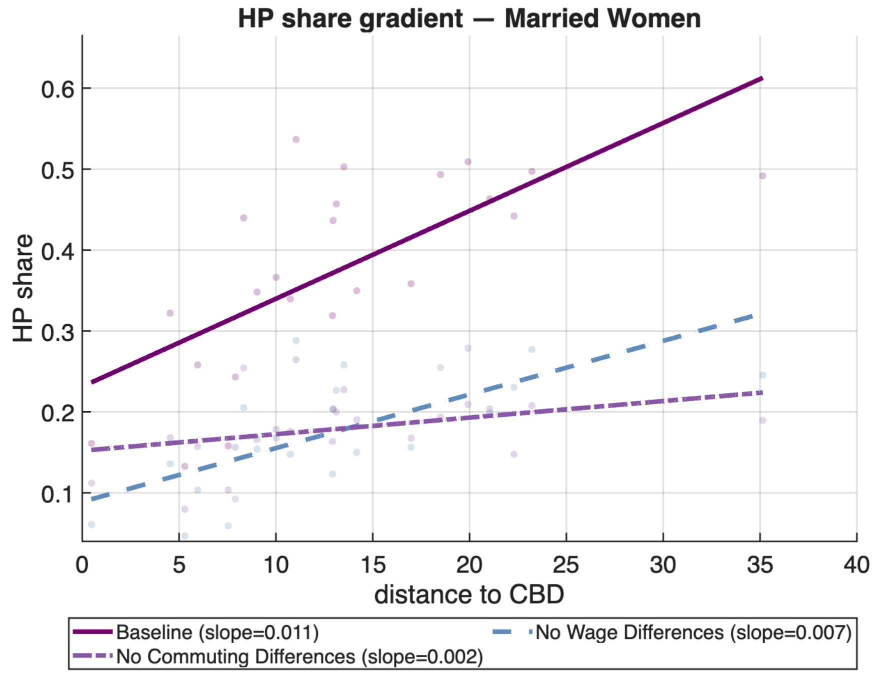

Notes: The figure shows the share of married women in home production by neighborhood against distance to the CBD (in kilometers). Each point represents a residential location, with fitted linear trends showing the gradient of home production with respect to distance from the CBD.

differences are removed, welfare rises by more than 17% for both single men and single women. The welfare increase for couples is more modest, but larger than when firm discrimination is removed. Equal access to the city's labor market reduces the relative disadvantage of women and peripheral residents, while also lowering commuting disutility for all workers. This shows that improvements in mobility can deliver broad welfare gains even without large changes in wages. The housing market plays a central role in mediating the aggregate effects of these policies. Figure 11 plots residential floor space prices by distance to the CBD. When wage differences are removed, higher wages concentrate demand in central neighborhoods, pushing up rents and steepening the housing-price gradient. The increases in floor space prices absorbs part of the welfare gains from higher wages. In contrast, equalizing commuting conditions flattens the rent gradient: as accessibility improves citywide, households are less constrained to pay a premium for centrality. The resulting equilibrium reallocates both labor and residence more efficiently across space, leading to smaller price pressures and higher real welfare as shown in Figure 7.

Figure 10: Counterfactual Changes in Welfare by Household Type

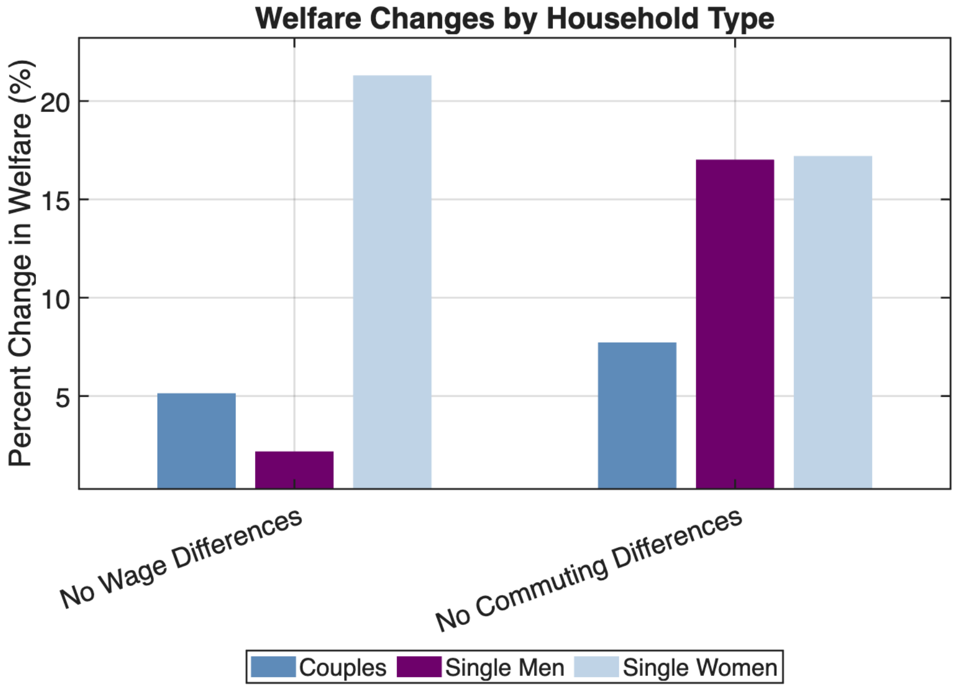

Notes: The bars show the percent changes in welfare relative to the baseline by household type. Blue bars correspond to couples, dark purple bars to single men, and light blue bars to single women.

Both experiments highlight distinct mechanisms linking gendered labor market frictions to the urban spatial equilibrium. Removing wage discrimination raises output by eliminating the taste-based disutility that firms experience when hiring women, which had previously reduced the share of women's production that entered aggregate output. However, the resulting rise in nominal incomes leads to higher housing prices and a mild redistribution of welfare: single women and couples gain, while welfare for single men declines slightly due to higher rents and limited wage adjustments. In contrast, removing commuting differences raises welfare more substantially and more uniformly across groups by improving effective access to jobs, reducing home production, and flattening spatial inequality in participation. Together, the counterfactuals reveal that mobility constraints are central to understanding gender gaps in labor force participation and the geography of home production within cities.

Figure 11: Counterfactual Residential Floor Space Prices

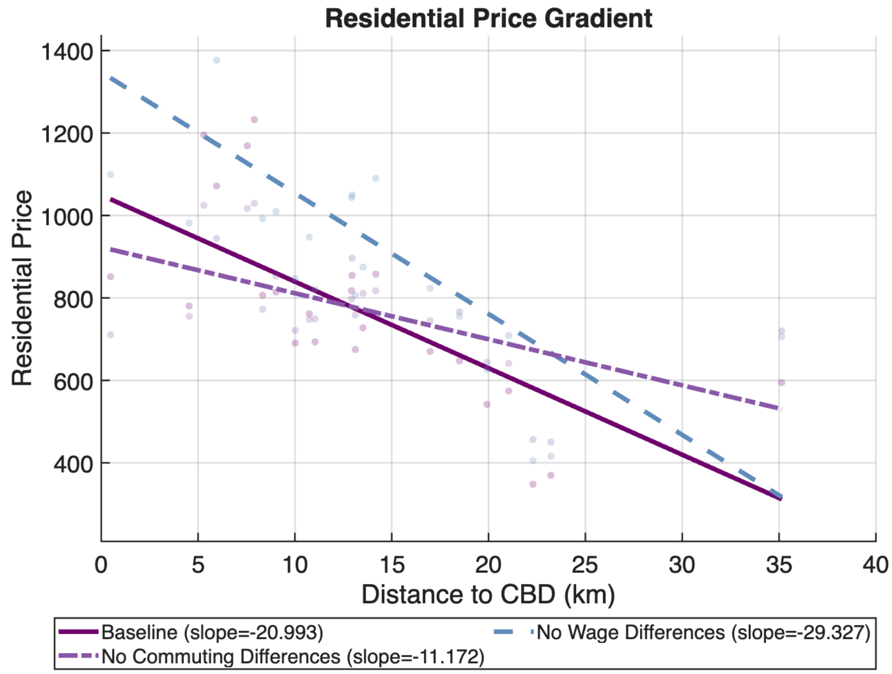

Notes: The figure shows residential floor space prices by distance to the CBD (in kilometers) for baseline and counterfactual scenarios. Each point represents a residential location, with fitted linear trends showing the price gradient.

6 Conclusion

This paper studies how commuting frictions shape labor supply within cities and contribute to gender gaps in market participation. Using rich microdata for São Paulo, I document that women, particularly married women, are less likely to work in the market sector when they face longer or slower commutes. These differences in commuting behavior are not simply a reflection of where people live, but a key driver of spatial and gender disparities in labor-force participation. To interpret these patterns, I develop a quantitative spatial model in which couples and singles jointly choose where to live and whether and where to work, trading off local amenities, housing costs, and commuting frictions. The model reproduces salient features of the data, including gender differences in commuting times, modes, and participation decisions.

Counterfactual exercises reveal that reducing gender gaps in commuting would generate large labor market and welfare gains. Equalizing commuting conditions between men and women raises female labor-force participation in the city by about ten percentage points, with the largest effects for married women living in the periphery. These changes translate into higher aggregate output and welfare, highlighting that gender differences in mobility are not only an issue of equity but also of efficiency. In cities where infrastructure and transport networks disproportionately disadvantage women, the economy forgoes a substantial share of its productive potential.

The results have clear policy implications. Urban transport and housing policies that improve accessibility for women, such as safer, more reliable, and better-integrated public transport or incentives for mixed-use development, can meaningfully increase women's participation in the labor market. Because commuting costs are shaped by both the built environment and household constraints, reducing spatial and gender inequalities requires coordination between urban and social policy. Measures that expand access to private or flexible modes of transport or support dual-earner households in peripheral neighborhoods are likely to have large returns in terms of both welfare and aggregate output.

Beyond the specific context of São Paulo, the paper contributes to a broader understanding of how urban structure interacts with household decisions. An important avenue for future research is to understand the origins of gender differences in commuting and identify which policies can most effectively reduce them. Differences in mode choice, access to private vehicles, safety concerns, and time constraints within households may all play a role in shaping mobility gaps. Understanding these mechanisms is essential to design targeted policies that promote women's access to jobs across urban areas.

References

Aguilar, A., Gutiérrez, E., and Villagrán, P. S. (2021). Benefits and unintended consequences of gender segregation in public transportation: Evidence from mexico city's subway system. *Economic Development and Cultural Change*, 69(4):1379–1410.

Ahlfeldt, G. M., Mulalic, I., Soto-Vieira, C., and Sturm, D. M. (2025). The geography of life: Evidence from copenhagen. *Working Paper*.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The Economics of Density: Evidence From the Berlin Wall. *Econometrica*, 83(6):2127–2189.

Alba-Vivar, F. (2024). Opportunity bound: Transport and access to college in a megacity. *Working Paper*.

Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3):1085–1140.

Bandiera, O., Elsayed, A., Smurra, A., and Zipfel, C. (2022). Young adults and labor markets in africa. *Journal of Economic Perspectives*, 36(1):81–100.

Baum-Snow, N., Gendron-Carrier, N., and Pavan, R. (2024). Local productivity spillovers. *American Economic Review*, 114(4):1030–1069.

Becker, G. S. (1965). A theory of the allocation of time. *The Economic Journal*, 75(299):493–517.

Becker, G. S. (2010). *The economics of discrimination*. University of Chicago press.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713.

Black, D. A., Kolesnikova, N., and Taylor, L. J. (2014). Why do so few women work in New York (and so many in Minneapolis)? Labor supply of married women across US cities. *Journal of Urban Economics*, 79:59–71.

Borghorst, M., Mulalic, I., and van Ommeren, J. (2024). Commuting, gender and children. *Journal of Urban Economics*, 144:103709.

Borker, G. et al. (2021). Safety first: Perceived risk of street harassment and educational choices of women. *World Bank Washington, DC*.

Buchmann, N., Meyer, C., and Sullivan, C. D. (2024). Paternalistic discrimination. *Stanford Institute for Economic Policy Research (SIEPR)*.

Bütikofer, A., Karadakic, R., and Willén, A. (2025). Parenthood and the gender gap in commuting. *Journal of Public Economics*, 248:105371.

Bütikofer, A., Løken, K. V., and Willén, A. (2024). Building bridges and widening gaps. *Review of Economics and Statistics*, 106(3):681–697.

Caldwell, S. and Danieli, O. (2024). Outside options in the labour market. *The Review of Economic Studies*, 91(6):3286–3315.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica*, 87(3):741–835.

Carta, F. and De Philippis, M. (2018). You've come a long way, baby. husbands' commuting time and family labour supply. *Regional Science and Urban Economics*, 69:25–37.

Cavalcanti, T. and Tavares, J. (2016). The output cost of gender discrimination: A model-based macroeconomics estimate. *The Economic Journal*, 126(590):109–134.

Cheema, A., Khwaja, A. I., Naseer, F., and Shapiro, J. N. (2019). Glass walls: Experimental evidence on access constraints faced by women.

Chen, Y., Coşar, K., Ghose, D., Mahendru, S., and Sekhri, S. (2024). Gender-specific transportation costs and female time use: Evidence from india's pink slip program. Technical report, National Bureau of Economic Research.

Christensen, P. and Osman, A. (2023). The demand for mobility: Evidence from an experiment with uber riders. Technical report, National Bureau of Economic Research.

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

Cuberes, D. and Teignier, M. (2016). Aggregate effects of gender gaps in the labor market: A quantitative estimate. *Journal of human capital*, 10(1):1–32.

Dasgupta, A. and Datta, A. (2023). Gendered transport subsidy and its short run effect on female employment: Evidence from delhi's pink pass scheme. Technical report, Working Paper.

Fan, J. and Zou, B. (2021). The dual local markets: Family, jobs, and the spatial distribution of skills. *SSRN Electronic Journal,(November)*.

Farré, L., Jofre-Monseny, J., and Torrecillas, M. (2023). Commuting time and the gender gap in labor market participation. *Journal of Economic Geography*, 23(4):847–870.

Ferreira, L. S. (2019). *Vegetacao, temperatura de superficie e morfologia urbana: um retrato da regiao metropolitana de So Paulo*. PhD thesis, Universidade de Sao Paulo.

Fiala, N., Garcia-Hernandez, A., Narula, K., and Prakash, N. (2022). Wheels of change: Transforming girls' lives with bicycles. *JSTOR*.

Franklin, S. (2020). Enabled to work: The impact of government housing on slum dwellers in south africa. *Journal of Urban Economics*, 118:103265.

Garlick, R., Field, E., and Vyborny, K. (2025). Women's mobility and labor supply: Experimental evidence from pakistan. IZA Discussion Paper 17883, IZA Institute of Labor Economics.

Glaeser, E. L. and Gottlieb, J. D. (2009). The wealth of cities: Agglomeration economies and spatial equilibrium in the united states. *Journal of Economic Literature*, 47(4):983–1028.

Harari, M. (2024). Residential patterns and local public goods in urban brazil. Technical report, Working paper.

Heath, R., Bernhardt, A., Borker, G., Fitzpatrick, A., Keats, A., McKelway, M., Menzel, A., Molina, T., and Sharma, G. (2024). Female labour force participation. *VoxDevLit*, 11(1):1–43.

Ho, L., Jalota, S., and Karandikar, A. (2024). Bringing work home: Flexible arrangements as gateway jobs for women in west bengal. *Structural Transformation and Economic Growth (STEG) Working Paper*, 80.

Hojman, A. and Boo, F. L. (2022). Public childcare benefits children and mothers: Evidence from a nationwide experiment in a developing country. *Journal of Public Economics*, 212:104686.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica*, 87(5):1439–1474.

Jalota, S. and Ho, L. (2024). What works for her? how work-from-home jobs affect female labor force participation in urban india. *Working Paper*.

Jayachandran, S. (2020). Social norms as a barrier to women's employment in developing countries. Technical report, National Bureau of Economic Research.

Joshi, S. (2024). Spatial shocks and gender employment gaps: Evidence from rising import competition in india. Working paper.

Juhn, C. and McCue, K. (2017). Specialization then and now: Marriage, children, and the gender earnings gap across cohorts. *Journal of Economic Perspectives*, 31(1):183–204.

Kawabata, M. and Abe, Y. (2018). Intra-metropolitan spatial patterns of female labor force participation and commute times in tokyo. *Regional Science and Urban Economics*, 68:291–303.

Kondylis, F., Legovini, A., Vyborny, K., Zwager, A., and Andrade, L. (2025). Demand for safe spaces: Avoiding harassment and complying with norms. *Journal of Development Economics*, 174:103392.

Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426.

Liao, Y., Gil, J., Pereira, R. H., Yeh, S., and Verendel, V. (2020). Disparities in travel times between car and transit: Spatiotemporal patterns in cities. *Scientific reports*, 10(1):4056.

Liu, S. et al. (2025). Why is female labor force participation declining in china? a perspective from urban commuting. *Journal of Development Economics*. Forthcoming/online first.

Liu, S. and Su, Y. (2022). The Geography of Jobs and the Gender Wage Gap. *The Review of Economics and Statistics*, pages 1–27.

Mincer, J. (1962). Labor force participation of married women: A study of labor supply. In *Aspects of Labor Economics*, pages 63–105. Princeton University Press for NBER, Princeton, NJ.

Monte, F., Redding, S. J., and Rossi-Hansberg, E. (2018). Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12):3855–3890.

Moreno-Maldonado, A. (2022). Mums and the city: Household labour supply and location choice. *Available at SSRN 3466171*.

Moreno-Maldonado, A. and Santamaria, C. (2022). Delayed childbearing and urban revival. *Working Paper*.

Ngai, L. R. and Petrongolo, B. (2017). Gender Gaps and the Rise of the Service Economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.

Oreffice, S. and Sansone, D. (2023). Commuting to work and gender norms by sexual orientation. *Labour Economics*, 85:102451.

Pereira, R. H., Saraiva, M., Herszenhut, D., Braga, C. K. V., and Conway, M. W. (2021). r5r: rapid realistic routing on multimodal transport networks with r 5 in r. *Findings*.

Petrongolo, B. and Ronchi, M. (2020). Gender gaps and the structure of local labor markets. *Labour Economics*, 64:101819.

Pietrabissa, G. (2023). School access and city structure. *Unpublished manuscript*.

Ranosova, T. (2025). Commuting and the value of marriage. *Working Paper*.

Redding, S. J. (2023). Quantitative urban models: from theory to data. *Journal of Economic Perspectives*, 37(2):75–98.

Redding, S. J. (2024). Quantitative Urban Economics. NBER Working Paper 33130, National Bureau of Economic Research.

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative Spatial Economics. *Annual Review of Economics*, 9(1):21–58.

Rosenthal, S. S. and Strange, W. C. (2004). Evidence on the nature and sources of agglomeration economies. In Henderson, J. V. and Thisse, J.-F., editors, *Handbook of Regional and Urban Economics, Vol. 4*, pages 2119–2171. Elsevier.

Seki, M. and Yamada, E. (2020). Heterogeneous effects of urban public transportation on employment by gender: Evidence from the delhi metro. *JICA Research Institute Working Paper*, (207):1–57.

Sharma, G. (2023). Monopsony and gender. *Unpublished manuscript*.

Silva, J. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and statistics*, pages 641–658.

Smits, J. and Steendijk, R. (2015). The international wealth index (iwi). *Social indicators research*, 122(1):65–85.

Tsivanidis, N. (2022). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogota's TransMilenio. *Working Paper*.

United Nations (2018). *World Urbanization Prospects: The 2018 Revision*. United Nations, New York.

Velásquez, D. (2022). Transit infrastructure and couples' commuting choices in general equilibrium. *Working Paper*.

White, M. J. (1986). Sex differences in urban commuting patterns. *The American economic review*, 76(2):368–372.

Wooldridge, J. M. (2015). Control Function Methods in Applied Econometrics. *Journal of Human Resources*, 50(2):420–445.

Zipfel, C. et al. (2022). *The demand side of Africa's demographic transition: desired fertility, wealth, and jobs*. STICERD, London School of Economics and Political Science.

A Data and Context

In this Appendix Section, I first provide some additional context about the city and where it stands relative to other large cities in the world. Then, I describe with detail the datasets used in the paper, and the creation of the relevant variables, providing additional description of the sample. Then, I discuss in detail how I created the travel time matrices used in the model estimation.

A.1 Additional Context

São Paulo city is part of the São Paulo Metropolitan Region, which comprises 39 municipalities in total. The region has a population of 19.6 million inhabitants, more than half (57.2 percent) of whom live in São Paulo municipality alone. It is the fourth most populous metropolitan region in the world. Although the city covers only 19.2 percent of the region's area, it concentrates 56.2 percent of its population. Nearly 85 percent of those who work in the city also reside there. Figure A1 highlights São Paulo within the metropolitan region and shows the distribution of Local Climate Zones, distinguishing areas of dense and sparse urban occupation.[26] Most of the densely built-up areas in the region are located within the municipality of São Paulo.

I focus on the city of São Paulo rather than the entire metropolitan region due to data limitations. Data on commercial floor space prices are unavailable for the other municipalities in the region. In addition, public transport networks are largely managed at the municipal level, which makes consistent network data across the region difficult to obtain. This is particularly binding for the model estimation, which depends on these additional variables. The stylized facts presented in Subsection 3.3 are robust to considering all the metropolitan region, rather than just the city.

A.2 Origin-Destination Surveys

The main data source used in this paper are the Origin-Destination Surveys for São Paulo. The *Pesquisa Origem e Destino* is a large-scale household survey conducted by the Deparment of Metropolitan Transportation of São Paulo Sate. It collects information on residents' daily travel patterns, socioeconomic characteristics, and mobility behavior. The 2007 wave covered around 30 thousand randomly selected households and 120 thousand individuals, distributed across 460 survey zones that span the entire São Paulo

[26] Ferreira (2019) classify satellite imagery according to the Local Climate Zone scheme.

Figure A1: São Paulo Metropolitan Region and Urban Occupation Type

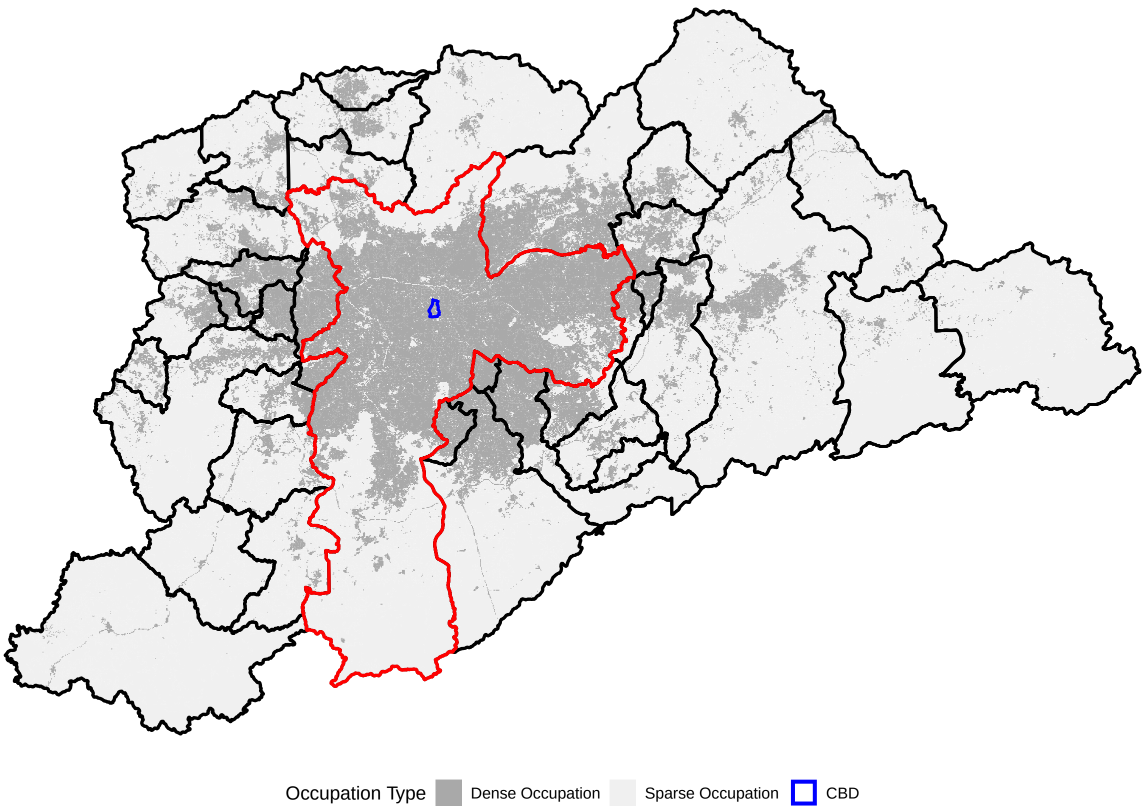

Notes: The map shows the 39 municipalities of the São Paulo metropolitan region and São Paulo municipality is highlighted in red. The data contains 100m2 satellite images classified into Local Climate Zones (LCZ). Dark grey considers any urban occupation, from compact high-rise (DN 1) until large lowrise (DN 8). The light grey area accounts for water, vegetation and sparsely built area (DN 9). Data and shapefile are from Ferreira (2019). Highlighted in blue is the center business district (CBD).

Metropolitan Region.

The survey provides micro data structured at the household, individual, and trip levels. The household questionnaire records information on dwelling characteristics, tenure, and ownership of durable goods and vehicles. At the individual level, it includes demographics, education, employment status, occupation, income, and schooling or workplace locations. The trip module reports all journeys made on the day prior to the interview, detailing origin and destination zones, purpose, departure and arrival times, modes of transport (up to four per trip), walking segments, and travel costs. This structure enables the construction of complete home–work commutes and the analysis of mode choice, travel time, and labor market accessibility.

Figure A2: Cities with Longest Commuting Time in the World

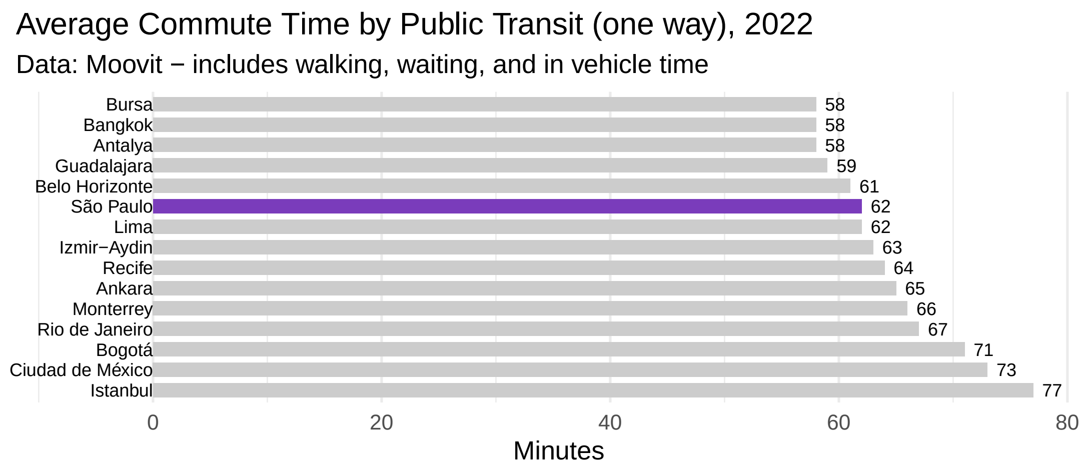

Definition of key variables. To identify couples, I use the household position variable reported in the survey. I classify as couples the household head and their reported spouse or partner, provided that both individuals are cohabiting in the same household. Each person is coded according to their relationship to the household head, which can be spouse or partner, child or stepchild, other relative, non-relative household member, live-in employee, or relative of an employee. Because this classification is defined relative to the household head, I am unable to detect additional couples within the same household when neither partner is the head. As a result, my measure may undercount cohabiting couples in extended or multi-family households. Over 65 percent of individuals in the sample live in households headed by a couple, while around 20 percent reside in non-nuclear or multi-generational arrangements that include additional adults beyond parents and children. These cases range from siblings or friends sharing a dwelling to extended families cohabiting in the same household. All remaining adults who are not classified as part of a couple are considered singles in the analysis.

I construct a measure of household wealth quintiles based on asset ownership, following the methodology commonly used in the Demographic and Health Surveys (Smits and Steendijk, 2015). The measure relies on information about durable goods and household amenities, under the assumption that variation in ownership and access to these assets captures long-run differences in economic status. Specifically, I include the number of bathrooms and domestic employees, as well as ownership of items such as cars, washing machines, refrigerators (one- or two-door), freezers, microwaves, dishwashers, DVD players, motorbikes, and bicycles. I apply Principal Component Analysis (PCA) to these asset variables and use the first principal component as a summary wealth index. Households are then ranked by their PCA score and divided into five equally sized groups (quintiles) using sampling weights. These measures are highly correlated with income and education level of the household head.

To define the sample for the analysis I use a variable called activity status, which is tabulated by gender in Table A1. I remove from the analysis students and retirees, who account for less than 8% of the adult sample aged between 18 to 55, and keep only individuals who are working or not working and house wives.

Table A1: Activity status by gender

Activity Status	Male (%)	Female (%)
Worker	84.17	63.86
House wife	0.18	17.07
Not working	10.38	13.38
Retired	2.63	3.2
Student	2.37	2.48
N	16638	19481

Source: Household Travel Survey of 2007. The sample contains individuals aged 18 to 55, residing in the municipality of São Paulo. Working also includes occasional workers, and any type of job or employment.

A.3 Travel time matrix

I compute the travel time matrices between neighborhoods by first identifying the population-weighted centroid of each neighborhood. Using a multimodal transport network router (Pereira et al., 2021), I calculate travel times between these centroids based on network data from OpenStreetMap and General Transit Feed Specification (GTFS) files for S ao Paulo. The network includes all major public transport modes—underground and overground rail, as well as municipal and inter-municipal buses—managed by two separate transport authorities.

In S ao Paulo, two agencies are responsible for maintaining GTFS data. *SPTrans*, the municipal transport management company, provides route information for buses and rail services operating within the city. The earliest available data from *SPTrans* are from January 2012. I adapt this network to represent conditions in 2007 by removing rail lines and stations built after that year and assuming limited changes to the bus network. The inter-municipal bus network, managed by the Metropolitan Urban Transport Department (EMTU), represents a smaller share of total transit. As the earliest EMTU data are from 2017, I use the same network to compute travel times for both 2007 and 2017.

The routing algorithm returns a distribution of travel times between each pair of centroids, reflecting variation in possible routes, departure times, and service frequencies. I calculate travel times for trips departing on a Friday at 6 a.m. and use the median travel time between each pair of points. Separate matrices are computed for public transport, private car, and walking. To obtain effective travel times by group, I weight these mode-specific matrices using observed mode shares by gender and marital status. Table A2 reports these shares, showing that women—especially married women—use public transport more intensively than men, while men rely more on private transport. Walking represents a small but non-negligible share across all groups, particularly among women. Public transport is the predominant commuting mode across all groups, indicating that most of the city's workers rely on it as their main means of travel to work.

Table A2: Transport Mode Shares by Group for Commuting Trips

	Public Transport	Private Transport	Walking
Married Women	0.58	0.23	0.19
Married Men	0.44	0.44	0.12
Single Women	0.65	0.19	0.16
Single Men	0.54	0.32	0.14

Note: Mode shares are computed from the 2007 Origin–Destination survey for employed individuals aged 18–55 living in the city of S ao Paulo. Private transport includes cars and motorcycles.

To validate the median travel times, I use the reported trip duration between neighborhoods from the Origin-Destination survey, and I plot them against the computed travel times using the routing algorithm.

Table A3: Commuting Outcomes by Gender and Marital Status

	Dependent Variable		
	Log Commuting Distance (1)	Log Commuting Time (2)	Share Driving (3)
Female	-0.04	0.10***	-0.13***
	(0.04)	(0.03)	(0.01)
Married	0.14***	0.06*	0.05***
	(0.04)	(0.03)	(0.01)
Female × Married	-0.26***	-0.15***	-0.03
	(0.06)	(0.04)	(0.02)
Observations	19,325	19,325	19,325
Residence Bin FE	✓	✓	✓
Education FE	✓	✓	✓
Wealth quintile FE	✓	✓	✓

Notes: Each column reports coefficients from a regression of the outcome on gender and marital status indicators and their interaction. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Controls include age, age squared, household size,number of 0-3 year-olds, 4-10 year olds and number of 75+ in the household in the household. Fixed effects for education level, for wealth quintile, and for location bin are included. Location bins are based on the distance to the CBD, and each bin includes 25% of the sample.

Table A4: Different Returns in the Labor Market by Gender and Marital Status

	Log(Earnings)
	(1)
Married Woman	-0.31***
	(0.02)
Single Woman	-0.32***
	(0.02)
Married Man	-0.11***
	(0.02)
Age	0.05***
	(0.005)
Age^2	-0.00***
	(0.00)
Log Time at Work	0.03*
	(0.02)
Observations	12,043
R^2	0.49
Sector fixed effects	✓
Occupation fixed effects	✓
Job Type fixed effects	✓
Education fixed effects	✓

Notes: The dependent variable is the log of the individual's wage. Controls include age, job characteristics, and education. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.