Gendered Use of Cities: Evidence from Jakarta*

Muskan Jain†

October 31, 2026

Click here for the latest version

Abstract

This paper examines gender differences in the use of urban space and how the provision of public space affects these differences. I create a scalable measure of street presence by gender from hand-labelled and machine-classified street imagery of Jakarta, and combine it with labour force and travel surveys. Women account for roughly one quarter of the adults observed on the street. An accounting exercise attributes about 70 per cent of the gender gap in street presence to differences in employment, hours and the location of work. Outside work, women make more trips than men but travel more locally. Comparing streets near and far from newly opened neighbourhood parks, I find that street presence rises by similar proportions for women and men, while women make up close to half of the adults photographed inside the spaces. Estimates from a nested discrete choice model of trip participation and destination choice indicate that women's participation responds more strongly to local access: 31 per cent of women's visits to new facilities represent additional trips, compared with 11 per cent for men. The results suggest that the gains from public space depend on residents' existing alternatives and on population density, and that such provision addresses only part of the gender gap in use of the city.

JEL Codes: C55, C81, H41, J16, J22, O18, R23, R53
Keywords: public space, gender, urban amenities, street presence, trip-making, Jakarta

*I am grateful to my advisors, Guy Michaels, Gharad Bryan, and Stephen Machin, for their guidance. I am thankful to participants of the Labour, Development, Trade and Urban seminars at LSE, Warwick PhD Conference for their feedback and comments. Rashmi Singh and Jayalakshmi Karthikeyan provided excellent research assistance for training the machine learning model and Martha Hindriyani provided excellent field assistance. Shinji Okazaki's assistance in data work is much appreciated. I thank Digital Data Divide for data collection. I gratefully acknowledge financial support from the International Growth Centre (IGC) Small Grants programme, LSE Saw Swee Hock Southeast Asia Centre (SEAC) PhD Research Support Fund, and CEP fund for early career independent research projects at the London School of Economics. For answering my questions and facilitating data access, I thank officials at BPS DKI Jakarta, the Jakarta provincial offices for women's empowerment and child protection (DPPAPP), spatial planning and land administration, and parks and urban forestry, together with Badan Pertanahan Nasional and Polda Metro Jaya, and staff at RPTRA sites across the city. I am also thankful to JICA for providing access to the travel survey data. All errors are my own. I declare no other financial relationships or conflicts of interest relevant to this research. I acknowledge the use of Claude and ChatGPT in writing to improve the clarity of the paper.

†Department of Economics, London School of Economics. Email: m.jain14@lse.ac.uk

1 Introduction

Cities concentrate jobs, services and amenities, allowing residents to undertake a wide range of market and non-market activities. Yet who benefits from this concentration depends on how individuals use the city: where they work, which destinations they visit and how they travel between them. Women and men differ substantially along these margins, but much less is known about how these differences translate into their use of urban space. This paper examines gender differences in the use of cities and how the provision and location of neighbourhood public spaces affect urban activity.

Gender differences in the use of cities can arise through several economic margins. Differences in employment rates, occupations and workplace locations affect where women and men undertake market work, while travel costs influence which destinations they choose and how frequently they travel. These costs may themselves differ by gender because of safety concerns, household responsibilities and restrictions on mobility. Existing research shows that women accept lower wages for shorter commutes (Le Barbanchon et al., 2021), choose safer journeys at the expense of educational quality (Borker, 2021), and increase job search when commuting options improve (Garlick et al., 2025). These studies establish the importance of spatial frictions for labour-market and educational outcomes. Employment and commuting, however, represent only part of urban activity. Shopping, accompanying children, accessing services and recreation also involve travel and the use of public space, and their spatial patterns may differ between women and men.

I examine gender differences in urban activity using two complementary measures: trip-making and street presence. Travel surveys record the frequency, purpose and destinations of journeys, allowing me to study where women and men travel for work and non-work activities. Repeated street imagery measures street presence, defined as the number and gender composition of people visible at different locations in the city. The two measures capture different dimensions of urban activity. Trip-making measures flows of journeys between locations, whereas street presence measures occupancy at a particular moment, reflecting both the frequency of travel and the time individuals remain visible at a location. Together, they allow me to examine gender differences in where people travel and where they are observed in public space, and how local amenity provision changes these patterns.

I study Jakarta, Indonesia, a large and densely populated city with substantial gender differences in employment and travel behaviour. Jakarta provides a useful setting for three reasons. First, residential areas, employment centres and public services are unevenly distributed across its neighbourhoods, making the location of destinations and the costs of reaching them relevant to how residents use the city. Second, repeated Google Street View imagery from 2013 to 2024 allows me to measure street presence across locations and over time, while household travel surveys provide information on trip frequency, purpose and destinations. Third, between 2015 and 2019, Jakarta introduced 315 *Ruang Publik Terpadu Ramah Anak* (RPTRA), or child-friendly integrated public spaces, across 166 of its 261 neighbourhoods. These are small, staffed neighbourhood facilities combining recreational areas with services for children, families and local communities. Variation in their locations and opening dates allows me to examine how new public spaces change street presence, while the travel data support a model of trip-making and counterfactual facility placement.

I measure the gender composition of people observed on the streets of Jakarta by creating two datasets using street imagery. The first consists of 10,216 hand-labelled Google Street View panoramas at 2,963 locations between 2013 and 2024. For each panorama, I record the gender and age group of people

present, their posture, activities and features of the surrounding built environment. This sample provides a direct measure of gender differences in street presence and allows me to characterise the activities undertaken by women and men in public space. The second dataset extends measurement to a substantially larger set of locations and years using computer vision. I develop a pipeline that detects pedestrians and estimates gender-specific probabilities using a classifier trained on publicly available pedestrian datasets and independently labelled images from Jakarta. I apply this pipeline to approximately 1.4 million images from $234{,}000$ panoramas collected between 2013 and 2019, constructing a panel of expected female and male pedestrian counts. I validate these measures against the separately hand-labelled street imagery and examine the sensitivity of the results to classification errors and uneven imagery coverage. I combine this data with data from labour force, commuter and travel surveys, which allows me to document gender differences in street presence and activities across the city and estimate changes in pedestrian presence around newly opened public-space facilities.

Using street imagery, I find that women account for approximately one-quarter of identified adults observed on Jakarta's streets, compared with about half of the city's adult population. The gender gap remains broadly stable across the observed years and daylight periods, but varies considerably across Jakarta's neighbourhoods. The share of women among adults observed on streets is lowest in areas with high employment density. Smaller samples of street imagery from other cities also reveal gender gaps, with women accounting for between 19.0 per cent of observed adults in Delhi and 43.2 per cent in New York.[1] These comparisons suggest that women's lower presence in public space is not unique to Jakarta, although the magnitude of the gap varies considerably across cities.

I first examine the contribution of work to the gender gap in street presence. Labour supply affects street occupancy through employment, working hours, the location and visibility of work, and commuting. Individuals working at home or inside buildings contribute differently to street presence from those working in street-facing occupations, even when their employment status and hours are similar. I combine labour-force and workplace information with commuting data to construct a measure of visible work-related time by gender. This accounting exercise relates the observed street gap to employment, hours, workplace visibility and commuting.

Under the baseline assignments of workplace visibility, employment accounts for 40.3 per cent of the street-presence gap and workplace visibility for a further 24.1 per cent. Differences in working hours account for a smaller component, while measured commuting partly offsets the gap. Together, these four margins account for at most 69.8 per cent of the observed difference under the baseline assumptions. This is a conditional bound obtained by assigning women zero visible time outside measured work. Its magnitude depends on the visibility assigned to occupations and workplaces. Employment and workplace location account for the largest components of measured work-related exposure, while the residual gap reflects differences in non-work activities and other sources of visible time.

I next compare street presence with travel behaviour using the 2002 SITRAMP household travel survey. Women make about 82 per cent as many trip legs per adult as men, compared with a female-to-male street-presence ratio of approximately 33 per cent. The travel survey and imagery cover different periods, so these ratios do not constitute a decomposition of the street gap. Their difference is consistent with gender variation in trip purposes, destinations, transport modes and the time individuals remain visible during their activities. The number of journeys and the stock of people observed on streets

[1] Appendix Table **??** reports female shares from street imagery across Bangkok, Delhi, Nairobi, London and New York.

therefore capture different aspects of urban activity.

The distribution of trips by purpose also differs by gender. Women make fewer trips overall but more non-work trips than men: approximately 0.65 non-work legs per adult per day, compared with 0.46 for men. Women are more likely to undertake non-work travel, while the number of such trips conditional on travelling is similar across genders. Their destinations for non-work activities are also closer to home. By contrast, gender differences in the distances travelled for work and school are much smaller. These patterns suggest that proximity plays a greater role in women's non-work destination choices, motivating an examination of how local amenities affect activity outside employment.

I study this question using Jakarta's *Ruang Publik Terpadu Ramah Anak* (RPTRA) programme. RPTRAs are small, staffed neighbourhood public spaces offering playgrounds, recreation, community facilities and services for children and families. The city registered 315 facilities opened between 2015 and 2019 across 166 neighbourhoods. I use variation in their opening dates using a spatial difference-in-differences ring design that compares changes in pedestrian counts close to each facility's boundary with changes on streets 750–$1{,}500$ metres away. Cell fixed effects account for persistent differences between streets, while facility-by-year effects absorb changes shared by streets around the same RPTRA in a given year. Identification requires comparable proportional trends across distance bands in the absence of an opening. I examine event-study estimates, alternative comparison rings and exposure to nearby facilities to assess the sensitivity of the results.

Following an RPTRA opening, expected female pedestrian counts within 25 metres of the facility boundary increase by 31.3 per cent and male counts by 23.1 per cent relative to the comparison ring. The estimates become smaller with distance from the boundary. Although the female point estimate is larger, the difference between the proportional responses is not statistically distinguishable from zero. The estimates indicate an increase in nearby female and male street presence, without establishing a change in the gender composition of surrounding street activity. The effects remain positive under alternative spatial comparisons, although their magnitudes depend on the comparison ring.

I also examine whether the gender composition of people inside RPTRAs differs from that observed on surrounding streets. Using photographs uploaded to Google Maps, I find that women account for 47.8 per cent of identified adults inside RPTRAs, compared with 26.1 per cent in machine-classified street imagery within 250 metres of the same facilities after opening. The comparison is descriptive, as uploaded photographs are not a representative sample of park visitors. The hand-labelled street imagery provides additional evidence on gender differences in activities. Women account for only 16.8 per cent of seated adults, compared with 27.9 per cent of adults recorded as on foot. Women are also more frequently observed buying goods and accompanying others, whereas men are more frequently observed selling or working as street vendors. These differences suggest that the composition of activity inside RPTRAs may differ from that on surrounding streets. For instance, visitors may spend much of their time inside the facilities, while street vendors and other workers remain visible nearby for longer periods. This is one possible explanation for why greater female representation inside RPTRAs coexists with statistically similar proportional street-presence responses for women and men outside of the parks.

The reduced-form estimates describe changes in street activity around the RPTRAs Jakarta built. They do not distinguish additional trips from visits redirected from other destinations, nor do they show where an equivalent quantity of public-space land would generate the largest increase in trip-making. I examine these questions using a two-stage model of destination choice and trip-making. Conditional on

undertaking an activity, residents choose among destinations according to their attractiveness and distance from home, allowing distance sensitivity to differ by gender. The resulting destination-choice inclusive value measures access to that activity and enters the probability of making a trip. I estimate destination attractiveness and distance sensitivity from survey origin–destination flows, and the relationship between access and trip-making from individual travel records. The framework draws on random-utility destination choice (McFadden, 1974) and gravity-based quantitative urban models (Ahlfeldt et al., 2015; Redding and Rossi-Hansberg, 2017; Miyauchi et al., 2025). The counterfactuals hold residences, employment and existing destinations fixed, allowing me to isolate changes in access and trip-making associated with new facility locations.

For the RPTRA-related activity composite, the estimated distance coefficient is 1.02 for women and 0.77 for men, consistent with greater distance sensitivity in women's destination choices. The coefficient relating log access to trip-making log odds is also larger for women, 0.317 compared with 0.108 for men. The activity composite combines recreation with private trips to educational and medical destinations, reflecting the range of services provided by RPTRAs. I use these estimated travel relationships to evaluate new facilities under alternative assumptions about their attractiveness, rather than fitting the model to the RPTRA opening effects. The access-response coefficient is estimated from differences across neighbourhoods conditional on household characteristics and sub-district effects. The counterfactual applies this relationship to changes in access generated by new facilities.

The model separates visits to newly introduced facilities into additional trips and trips redirected from existing destinations. Under the main assumptions, the network corresponding to Jakarta's registered RPTRAs generates approximately $4{,}986$ additional female trips and $1{,}151$ additional male trips per day. Of predicted visits to the added facilities, 31 per cent for women and 11 per cent for men represent additional trips, with the remainder replacing trips to existing destinations. The larger predicted female increase is driven primarily by the estimated relationship between access and trip-making rather than the difference in distance sensitivity. Predicted visits to a facility therefore exceed the net increase in trips for the activity it provides.

I next compare alternative locations for a fixed quantity of public-space land. Holding total land at 55.78 hectares and allocating it across 100 equally sized facilities, I compare placement rules based on residential population density, low existing access and predicted total additional trips. Population-density targeting generates $4{,}510$ additional trips per day across women and men; prioritising low existing access generates $7{,}396$; and the best-found model-based allocation generates $7{,}874$. The latter represents approximately 75 per cent more additional trips than density targeting under the same facility-number and land constraints. These differences depend on the population reached, distances to new facilities and access to existing destinations. Allocations that separately maximise women's and men's additional trips also produce similar predicted female gains: choosing locations under the male objective lowers women's additional trips by 0.38 per cent relative to the female objective.

The empirical and model results measure related but distinct outcomes. The opening estimates concern the number of pedestrians visible around RPTRAs, while the interior photographs describe the gender composition of selected facility users. The model predicts changes in trip-making and distinguishes additional trips from substitution across destinations. Consequently, an increase in nearby pedestrian activity need not correspond to an equally large increase in total trip-making, and the gender composition inside a facility need not match that on surrounding streets. The results suggest that increasing local street

activity, generating additional trips and narrowing gender differences in public-space use are distinct policy objectives. The allocation counterfactual measures additional trips under a fixed land constraint, rather than monetary welfare or the full costs of facility provision.

The questions examined here are relevant to other cities where employment, services and amenities are unevenly distributed and where travel costs or activity patterns differ by gender. The smaller imagery samples from Delhi and New York suggest substantial variation in female representation across cities, although they do not establish the external validity of Jakarta's policy effects. Comparable street imagery and travel data could be used to examine whether gender differences in public-space occupancy are similarly related to work, travel and local amenity provision in other settings.

The paper contributes to three strands of literature. First, it adds a spatial outcome to research on gender differences in labour supply, occupational allocation and the division of time between market and non-market activities (Goldin, 2014; Kleven and Landais, 2017; Olivetti et al., 2024). Existing research examines employment, hours, occupational sorting and household production. I examine where the combined consequences of these activities become visible in the city. The accounting exercise connects street presence to employment, hours, workplace visibility and commuting, while the travel analysis documents gender differences in non-work destination choices. This complements research on commuting costs and gender-specific travel constraints (Le Barbanchon et al., 2021; Borker, 2021; Garlick et al., 2025) by studying the distribution of people across public space alongside labour-market and travel outcomes.

Second, the paper relates to urban economics on spatial frictions, destination choice and the provision of local amenities. Quantitative urban models use travel costs and spatial choice to study residential locations, employment and mobility (Ahlfeldt et al., 2015; Redding and Rossi-Hansberg, 2017; Miyauchi et al., 2025). I study changes in pedestrian occupancy following the introduction of neighbourhood public spaces and use a destination-choice model to compare additional trip-making under alternative facility locations. The reduced-form estimates concern local activity around actual openings, while the model distinguishes additional trips from substitution among destinations. The allocation exercise incorporates residential population, distance and existing access when comparing the predicted effects of new public-space provision.

Third, the paper contributes to the measurement of urban economic activity using spatial data. Satellite imagery has been used to measure economic conditions and land use (Donaldson and Storeygard, 2016), mobile-phone records to study travel and spatial interaction (Kreindler and Miyauchi, 2023; Miyauchi et al., 2025), and street imagery to characterise neighbourhoods and the built environment (Gebru et al., 2017; Glaeser et al., 2018; Harari and Wong, 2026). I use street-level photographs to measure the number and demographic composition of people present in public space. The approach complements camera-based street observation, including Karekurve-Ramachandra and Sood (2026), while using Street View's historical and geographic coverage to examine changes around local amenities. Repeated measurement of street presence also provides data for studying how transport, safety, commerce and other urban policies affect the spatial distribution of visible activity.

The remainder of the paper is organised as follows. Section 2 describes Jakarta, the RPTRA programme, and the data. Section 3 introduces the street-presence measure, while Section 4 and Section 5 examine work-related exposure and travel patterns. Section 6 studies RPTRA openings and activity inside and around the facilities. Section 7 presents the destination-choice and trip-making model and the

allocation counterfactuals. Section 8 concludes.

2 Setting and Data

2.1 Jakarta

DKI Jakarta, Indonesia's national capital, has provincial status and comprises five administrative cities (*regencies*) and the Thousand Islands regency, divided into 44 sub-districts (*kecamatans*) and 267 neighbourhoods (*kelurahans*). I exclude the offshore Thousand Islands and focus on the 261 mainland neighbourhoods, the smallest administrative units with consistent survey records. With 10.6 million residents at the 2020 census (Badan Pusat Statistik Provinsi DKI Jakarta, 2021), DKI Jakarta forms the core of Jabodetabek, the wider metropolitan region encompassing Bogor, Depok, Tangerang and Bekasi (Figure 1, panel a). Historical Street View coverage begins in 2013, providing observations from two years before the provincial government introduced neighbourhood public spaces through the RPTRA programme.

Within Jakarta, employment is more spatially concentrated than residence. Formal jobs cluster in the central business district, which extends south from the National Monument along Jalan Sudirman, while residential neighbourhoods are more widely distributed (Badan Pusat Statistik, 2017). This geography creates differences in the distances residents travel to work, schools, shops and recreational spaces, and brings workers and visitors into neighbourhoods other than their own. The people observed on a street consequently reflect the activities taking place there as well as the population living nearby.

Work contributes to these patterns differently for women and men, both through employment and through the location of their jobs. In the 2018 National Labour Force Survey (SAKERNAS), women's employment rate was about 62 per cent of men's, although the difference in usual hours among those employed was much smaller. The 2023 survey also records pronounced differences in workplaces: 6.4 per cent of employed women worked on streets, at roadsides or in itinerant locations, compared with 16.3 per cent of men, while women were more likely to work at home.

Beyond employment, household norms and safety concerns may affect where women and men undertake activities outside the home. In the 2017 Demographic and Health Survey, 15.5 per cent of urban married women and 6.5 per cent of urban married men reported that a husband was justified in beating his wife if she went out without telling him (ICF International, 2018). Concerns about personal safety were also more common among women: in the 2023 World Risk Poll, 82.4 per cent of women and 71.5 per cent of men in Indonesia expressed concern that violent crime could cause them serious harm (Lloyd's Register Foundation, 2023). Such concerns may influence the route, mode or timing of a journey, or the destinations considered acceptable, even when they do not prevent an activity altogether. Appendix A.11 compares these measures across countries.

These differential constraints and/or preferences can extend across work and non-work activity. Restrictions on independent travel may affect employment as well as leisure, while care responsibilities can shape both working hours and the journeys made outside work. Time-use and well-being surveys provide further context through gender differences in committed time, leisure activities and concern about travelling alone (Appendix A.10). I draw on labour-market and travel data to examine where women and men work and undertake other activities, relating these patterns to their observed street presence rather than estimating the separate contribution of each underlying constraint.

2.2 The RPTRA Programme

In 2015, Jakarta began expanding its network of neighbourhood destinations through the *Ruang Publik Terpadu Ramah Anak* (RPTRA), or Child-Friendly Integrated Public Space, programme. Introduced under Governor Basuki Tjahaja Purnama, the programme sought to bring public space, recreation and services for children and families closer to residential communities (Pemerintah Provinsi DKI Jakarta, 2015, 2016). RPTRAs are relatively small, formally managed spaces. Their standard facilities include playgrounds, sports areas, seating and greenery, alongside a multipurpose room, library, toilets and lactation room (Pemerintah Provinsi DKI Jakarta, 2016). Most are fenced, lit and monitored, with day-to-day management shared between the *kelurahan* and designated RPTRA staff.

The staff also organise early-childhood, community-health, exercise and children's activities, including those associated with the Family Welfare programme (PKK). Although the facilities share a common design, their services vary with the phase of construction, corporate social-responsibility (CSR) contributions, staff initiatives and local partnerships.[2] An RPTRA opening thus brings together outdoor and indoor space, staff and organised activities. The estimated effects concern this combination, which may attract women and men differently, rather than any individual feature of the facility.

Staffing and programming also make RPTRAs a recurrent public service rather than a one-off investment in land and construction. Available project and budget records put construction expenditure, excluding land, at roughly Rp0.5–0.7 billion for CSR-funded sites and Rp1.2–1.7 billion for those financed through the provincial budget. Records for 2017–2018 imply annual expenditure of about Rp0.4–0.45 billion per facility on staff, utilities, operations and maintenance.

The programme expanded quickly. The register contains 315 sites in 166 of Jakarta's 261 mainland neighbourhoods, with a combined area of 55.78 hectares and a median size of 1,559 square metres (Pemerintah Provinsi DKI Jakarta, 2021). The stated objective was to establish at least one RPTRA in each neighbourhood, and 270 sites (85.7 per cent) opened between 2015 and 2017, with the remainder opening by 2021. Figure 1, panel (b), shows their distribution; Appendix D reports opening dates and the number of facilities by neighbourhood. The sites opening after 2019 remain untreated during the street-imagery panel, while the model in Section 7 takes the full registered area as its land budget.

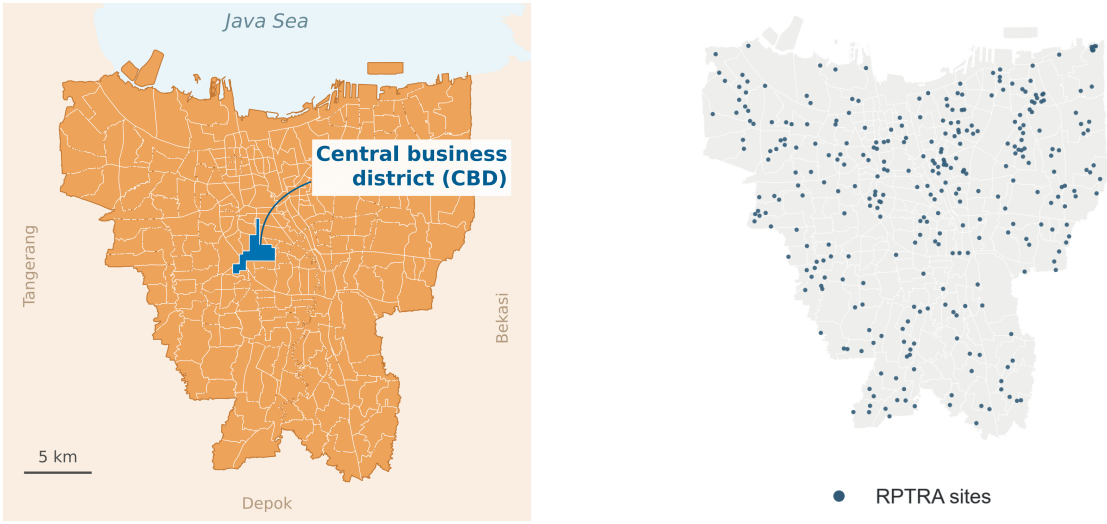
Available public land played an important role in determining where RPTRAs were built, alongside government decisions, provincial funding and corporate contributions (Berita Jakarta, 2016). Many were built on land already under public ownership or use. The programme was therefore not randomly located across Jakarta. In the main analysis, I compare changes in street presence near each site with changes farther from that same site, before and after it opens. The street outcome includes people passing a facility, arriving or leaving, and spending time near its boundary; it need not capture activity inside. I examine the composition of photographed users inside RPTRAs separately from changes on the surrounding streets.

2.3 Data

I study the use of the city through two complementary measures: street presence and trips. Street presence records who is visible in public space at a particular moment, while trips record whether people

[2] Some RPTRAs later acquired additional functions. Selected sites serve as bases for POS SAPA, local points for information, reporting and referral on violence against women and children, and some have received child-friendly-playground certification. These later designations are selected, and I do not treat them as part of the original intervention.

Figure 1: Jakarta neighbourhoods and the location of RPTRA sites



(a) Jakarta neighbourhoods and the CBD

(b) RPTRA sites

Notes: Panel (a) reproduces the slide map of mainland Jakarta neighbourhoods, highlighting the central business district (CBD) and labelling the surrounding areas. The analysis covers the 261 mainland neighbourhoods and excludes the Thousand Islands. The travel survey records journeys across the wider region; Section 2.3 describes the residential and trip restrictions. Panel (b) maps the 315 registered RPTRA sites against mainland neighbourhood boundaries, including sites that opened through 2021. It shows RPTRA sites only, not all public-space land.

travel, where they go and how they make their journeys. People who spend longer in public view are more likely to appear in the images, making repeated street observations informative about visible person-time; the diaries describe movement between places. I combine these measures with labour-market surveys and spatial data to examine how work, residential location and the availability of destinations shape urban activity.

Street presence. I measure street presence using Google Street View panoramas, which record the location and capture month and year of each image (Google, 2025b). I built two samples using street imagery. The first is a hand-labelled sample that contains 10,216 panoramas at 2,963 locations, captured between 2013 and 2024, excluding 2020–2021. Annotators recorded the apparent gender and life stage of each visible person, as well as posture, companionship and selected activities. I count pedestrians, excluding people in or on vehicles. These observations establish the level of the gender gap and describe how women and men are observed using the street.

To scale up the sample to allow for temporal and spatial comparisons, I built a larger machine-classified sample covering 2013–2019. It contains 233,933 panoramas, each processed in six directions, giving 1,403,598 images from repeated captures at approximately 65,000 locations. The classifier assigns probabilities of being a woman, man or child to detected pedestrians, which I sum to construct expected counts. Section 3 describes the sampling, classification, weighting and validation procedures.

Trips. The second measure comes from the 2002 JICA SITRAMP household travel survey (Japan International Cooperation Agency, 2002). It is a 3% random sample of households in Jakarta. Its daily diaries record each household member's journeys, including trip purpose, origin, destination, facility type, transport mode and travel time. They allow me to compare total trip-making by gender, as well as examine the purposes of those journeys, the distances travelled, and the time spent travelling. The sample

contains 185,744 adults aged 15–110 living in mainland Jakarta, comprising 92,014 women and 93,730 men. Importantly, the sample retains adults with valid diaries who made no trips, allowing me to compare travel across the full adult population rather than only among those who travelled. Unlike the commuter surveys, SITRAMP covers travel for non-work trips and records complete individual diaries, including days without travel.

The survey records trips for all purposes, allowing me to distinguish work, business, school, shopping, private and other journeys. I classify shopping, private and other trips as non-work activity, excluding return-home legs. These categories include errands and personal activities as well as recreation. Since one person may undertake several activities in a day, the purpose-specific measures are not mutually exclusive. I also distinguish individual trip legs from completed journeys that start and end at home, including return travel, when comparing journey durations.

I measure trip distances using the geographic locations of origins and destinations reported in the travel survey, which I map to Jakarta's 261 mainland neighbourhoods (*kelurahan*). For trips between different neighbourhoods, I calculate the straight-line distance between their centroids. For trips within the same neighbourhood, where origin and destination would otherwise have zero distance, I approximate the expected distance between two randomly located points in a circle with the same area as the neighbourhood. These measures provide a consistent approximation of spatial separation across trips, rather than the actual distance travelled along the road network. Appendix B describes the geographic matching and distance construction.

The survey predates the first Street View images by 11 years and the introduction of the RPTRA programme by 13 years. Although Jakarta's transport network and distribution of destinations have changed considerably since 2002, later surveys do not provide the comprehensive origin–destination flows needed to estimate the model. I use this data to examine gender differences in travel patterns and to estimate the destination-choice relationships used in the model, while recognising that travel behaviour may have changed over time.[3]

Work and commuting. Street presence includes working time in public view, which need not involve a journey. To relate the street counts to market work, I use the August 2018 and 2023 rounds of Labour Force Surveys (SAKERNAS), covering employment, usual working hours, and occupation among DKI Jakarta residents. The 2023 round records workplace location, distinguishing work at home, at another person's home, in offices and factories, in malls and covered markets, on streets and roadsides, and in itinerant locations. Combining workplace and occupation categories, I assign the fraction of working time likely to be visible from the street. The baseline accounting exercise uses employment and hours from 2018 and work locations from 2023, assuming that the later workplace distribution also applies to the earlier year.

Commuter Surveys by the National Statistics Agency (BPS) from 2014, 2019 and 2023 record journey modes and durations (Badan Pusat Statistik, 2014, 2019, 2023). Unlike SITRAMP, they cover only commutes crossing a regency boundary and therefore represent a selected set of journeys. I use these surveys to estimate visible commuting time. Section 4 and Appendix C.4 describe the accounting inputs in further detail.

[3]To assess the relevance of these historical patterns, I compare them with the 2018 JICA travel survey report (the microdata is not yet available for researchers). Although the magnitudes of travel outcomes differ across the two waves, the main gendered patterns in trip-making and destination choice remain similar.

Population and destinations. The relationship between travel and street presence also depends on where people live, work and find destinations. For the spatial comparisons and model, I use adult population counts by gender for the 261 mainland neighbourhoods from the ten-per-cent sample of the 2010 Population Census (Badan Pusat Statistik, 2010). The 2016 Economic Census supplies employment counts by neighbourhood (Badan Pusat Statistik, 2017).

The RPTRA register provides site names, opening dates, areas and locations. I digitise the boundaries of 298 facilities using Google Earth and approximate those of the remaining 17 using their registered areas and geo-locations. The register assigns the 315 sites to 166 neighbourhoods. Village Potential Statistics (PODES), administrative spatial files, Google Places and OpenStreetMap provide further neighbourhood and amenity characteristics (Badan Pusat Statistik, 2011–2024; Google, 2025a; OpenStreetMap contributors, 2025).

For the supplementary analysis of land values, I use assessed land-value maps (ZNT) prepared by the National Land Agency (BPN) for 2015 and 2019 (Kementerian Agraria dan Tata Ruang / Badan Pertanahan Nasional, 2015, 2019). The 2019 maps cover only North and South Jakarta, and the assessments are not transaction prices (Appendix D.9). I also use administrative population records compiled by Jakarta's Population and Civil Registration Agency (*Dinas Kependudukan dan Pencatatan Sipil, Dukcapil*) and made available through the Jakarta Satu Data portal. These records provide annual de jure population counts by gender and age group at the kelurahan level. Unlike census counts, they reflect registered residence, which may differ from where individuals actually live (Appendix D.8).

Activity inside RPTRAs. Street View generally does not show activity inside RPTRAs. I therefore hand-label 343 user-contributed Google Maps photographs from 208 facilities, distinguishing photographs showing ordinary use, organised events and posed groups. Of these, 217 photographs from 161 RPTRAs depict ordinary use. Section 6.3 compares the gender composition of photographed users with that of people on nearby streets. Because users choose what to photograph and upload, the images describe photographed composition rather than representative visitor attendance.[4]

Table A1 summarises the datasets, their reference periods and purpose of use in the analysis.

3 Measuring Street Presence

3.1 Street Presence and Journeys

Streets are both routes between destinations and places where people work, trade, wait and socialise. A single image may capture someone passing through, a worker at a roadside stall and a person waiting outside a shop. I measure this activity by counting adults visible on foot, standing or seated, excluding people inside moving motorised vehicles. The hand labels distinguish non-motorised riders, children and adults whose gender cannot be identified, allowing me to exclude these groups from the main pedestrian measure. Gender is assigned from visible appearance under the labelling protocol rather than reported by those photographed.

Street presence differs from the number of trips people make because it also reflects how long they

[4] In April 2026, I also visited nine RPTRAs, held semi-structured interviews with users and staff, and revisited three Street View locations at different times of day. These observations inform the discussion of possible mechanisms in Section 6.2 and the checks in Appendix A, but do not enter the estimates.

spend in public view. Someone passing through a street briefly is less likely to appear in an image than someone spending several hours working at a roadside stall. A person may make several journeys while remaining largely inside vehicles or buildings, whereas someone making fewer trips may spend much of the day working or socialising outdoors. Street presence therefore depends on both how often people travel and how long they remain in public view.

Repeated images capture this occupancy across locations and dates. Under comparable spatial and temporal capture, the ratio of female to male counts is informative about relative visible person-time during the periods photographed. Unlike travel surveys, the images do not rely on respondents recalling their journeys, but neither do they follow individuals or reveal how long each person remains in view. They also capture only part of urban activity: people inside shops, offices and community rooms are largely unseen, while waiting outside contributes to street presence without adding another trip.

To examine which activities contribute to this visible time, I distinguish between exposure associated with market work and commuting and exposure outside measured work. For adults of gender g, I write

$$V_g = W_g + \ell_g,$$

where V_g is visible person-time per adult, W_g is the work-related component constructed from employment, working hours, workplace visibility and commuting, and ℓ_g is visible time outside these measured activities. Section 4 develops the accounting using labour-market and commuter surveys.

The distinction between trips and street presence is substantial in Jakarta. Women make 82 per cent as many trip legs per adult as men in the 2002 household travel diaries, compared with a female-to-male pedestrian ratio of approximately 0.33 in the imagery (Table 2). Although the sources cover different periods, the gender gap in trip frequency is considerably smaller than the gap in street presence. Differences in travel modes, routes and time spent in public view may contribute to this contrast, alongside work and other activities undertaken on the street. I examine gender differences in the purpose and geography of journeys in Section 5.

3.2 Collecting the Images

To select locations for the hand-labelled sample, I first divided mainland Jakarta into a grid of 100-metre-square cells, giving 63,136 potential sampling locations. I then randomly selected 10 per cent of these cells (6,314) and searched for Google Street View imagery near the centre of each selected cell. Where imagery was unavailable at the centre, enumerators used the nearest eligible panorama within the same block. Locations without suitable imagery were excluded, as were privately uploaded images whose capture conditions differed from Google's vehicle-mounted imagery. Subject to a fixed collection quota and imagery availability, the final sample covered 2,963 locations, or 4.7 per cent of the original grid.

For each retained location, I collected images from multiple years where available. Enumerators viewed each panorama without zooming or navigating to other images to maintain consistent annotations. When a location has more than two historical captures, I retain the earliest and latest images, together with intermediate captures at least 24 months apart. The final hand-labelled sample comprises 10,216 panoramas from 2,963 locations, averaging 3.45 observations per location between 2013 and 2024, excluding 2020–2021.

To obtain repeated observations across a much larger part of the city, I construct a second sample

using a finer grid of hexagonal cells, each 50 metres wide. I oversample cells near RPTRA facilities to support the analysis of their openings. For each sampled cell and year in the study period, I search for Street View imagery within 25 metres of the cell centre and retain the nearest available panorama. I then download six views of the same panorama, each facing a different direction at 60-degree intervals, to capture the full surroundings. Between 2013 and 2019, this procedure yields 233,933 distinct panoramas and 1,403,598 directional images from 66,459 locations. I classify the people detected in these images and aggregate their gender probabilities into the counts used in the analysis.

Street View does not cover all roads in Jakarta. Usable imagery is unavailable for 57.1 per cent of sampled cells. Coverage is particularly uneven around RPTRAs: approximately 52 per cent of sampled cells within 100 metres of a facility have imagery, compared with about 20 per cent of cells more than two kilometres away. Among cells with imagery, however, the number of available captures varies much less with distance. Figure A1 illustrates the coverage problem in Sunter, North Jakarta, where Street View covers the main roads but misses many narrower residential lanes. I account for differences in sampling and imagery availability by weighting observations by the inverse of their sampling probabilities and estimated coverage probabilities.

Although Street View records the month and year of each panorama, it does not provide the exact date or time of capture. I therefore cannot distinguish weekday from weekend activity or observe how street presence changes over a representative day. For the hand-labelled sample, I use the direction of shadows to classify images into broad daylight periods wherever possible and examine whether the gender gap varies across them. The panel regressions also include capture-month fixed effects. All observations are restricted to daylight, and the images do not allow me to identify or follow individuals across locations or dates.

3.3 Hand Labelling

Digital Data Divide, a data-labelling firm, collected and annotated panoramas under a common instrument. Enumerators counted people visible on roads, pavements, shaded stalls and open seating, excluding private terraces, balconies and occupants of moving motorised vehicles. They distinguished infants, children and adults, and coded adults as women, men or unidentifiable. For each person, they also recorded posture (sitting/on foot), companionship (alone or in groups), selected activities (including buying, selling, and street vending), and whether women were wearing head coverings. Separate fields described the road, pavement, vehicles, greenery, seating, lighting and congestion. Appendix A.1 contains the coding rules.

Summing across the hand-labelled panoramas gives 6,029 identified women and 18,266 identified men on foot or seated. Women therefore account for 24.8 per cent of identified adult pedestrians.[5] Gender cannot be determined for approximately 8.7 per cent of the adult count, often because people are distant or obscured. I exclude them from the identified-adult denominator and examine alternative assignments to women or men in the sensitivity analysis.

[5] Including non-motorised riders, who are coded separately, yields 24.7 per cent. I use the pedestrian definition for the main analysis.

3.4 Training and Calibrating the Classifier

To construct gender-specific pedestrian counts from the larger Street View sample, I use a two-stage machine-learning pipeline. First, a YOLO object detector identifies people in each image and draws a bounding box around them.[6] I then crop each detected person and use a separate classifier to predict whether they are a woman, man or child. Detecting and classifying people is particularly challenging when they are distant, partially obscured or in crowded scenes.

I train the attribute classifier in two stages. I first use two publicly available pedestrian datasets: PA-100K, containing 100,000 images from 598 outdoor cameras, and PETA, containing 19,000 images from ten surveillance systems (Liu et al., 2017; Deng et al., 2014). I then fine-tune the classifier on independently labelled pedestrian images from Jakarta, allowing it to learn from the people and street conditions in the study setting. The Jakarta annotations distinguish women, men, children and unidentified people and are divided into training, validation and test sets in proportions of 70, 15 and 15 per cent, respectively.

For each detected person, the classifier produces a score for each category: woman, man and child. I use Dirichlet calibration, fitted on a separate human-labelled sample, to adjust these scores so that the resulting probabilities better reflect the observed frequencies of each category (Kull et al., 2019). I then calculate expected pedestrian counts by summing the calibrated probabilities across retained detections.

$$\widehat{N}_{\text{women}} = \sum_{r} \widehat{p}_r(\text{woman}).$$

where r is a detected individual. For example, a person assigned a 70 per cent probability of being a woman contributes 0.7 to the female count. I construct the male and child counts in the same way.

Figure 2 illustrates the detector's bounding boxes and the calibrated probability assigned to the woman category for each detected person.

I evaluate the calibrated probabilities using the multiclass Brier score, which measures the average squared difference between predicted probabilities and observed categories (Brier, 1950). In weighted cross-validation on the audited crops, calibration reduces the score from 0.460 to 0.278, indicating more accurate probability predictions. The Brier score evaluates classification accuracy among detected people; missed pedestrians and ineligible detections can still affect the final counts.

3.5 Validation

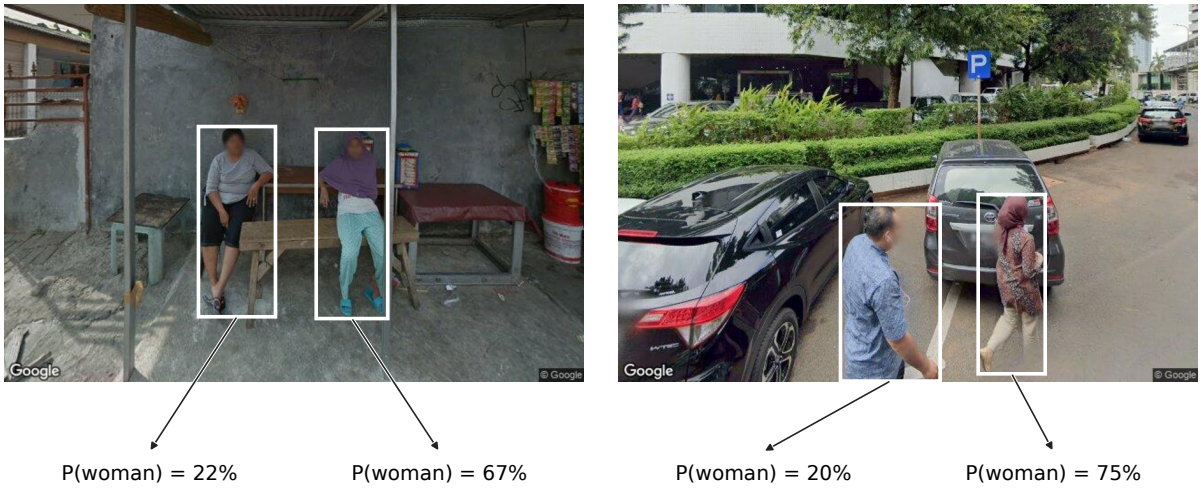
I assess the machine-classified measure using hand-labelled observations, errors associated with the RPTRA treatment design and independent mobile-phone mobility data (Table 1).

Comparison with hand-labelled data. Restricting the hand-labelled sample to 2013–2019, women account for 25.6 per cent of identified adults,[7] compared with 23.8 per cent in the machine-classified sample. The geographic distribution is also similar across the two samples. Both show higher female representation in East Jakarta and lower representation in Central Jakarta, with broadly similar shares across the five administrative cities (Figure 3, panel a).

[6] YOLO (You Only Look Once) is a computer-vision algorithm that detects objects in an image and identifies their locations using bounding boxes.

[7] The headline female pedestrian share of 24.8 per cent is calculated from the full hand-labelled sample, including later years.

Figure 2: From detected people to probabilistic counts



(a) Example 1 **(b)** Example 2

Notes: White boxes locate detected people in two Google Street View images. Labels report the Dirichlet-calibrated probability of the woman category, rounded to whole percentages; they are not detector-confidence scores or hand labels. The classifier assigns probabilities to woman, man and child, which sum to one for each detection; only the woman probability is displayed. Summing the displayed probabilities gives approximately 0.89 and 0.95 expected women in panels (a) and (b), respectively. These images were selected for readable illustration and do not constitute a random audit or an estimate of accuracy.

The classification audit shows that the model performs better for men than for women. Among hand-audited detections, precision, the proportion of predicted classifications that are correct, is 79.7 per cent for women and 85.1 per cent for men. Recall, the proportion of correctly identified cases within each true category, is 76.0 per cent for women and 90.4 per cent for men. The corresponding F1 scores, which combine precision and recall, are 77.8 and 87.7 per cent for women and men respectively. These statistics measure classification performance conditional on successful pedestrian detection.

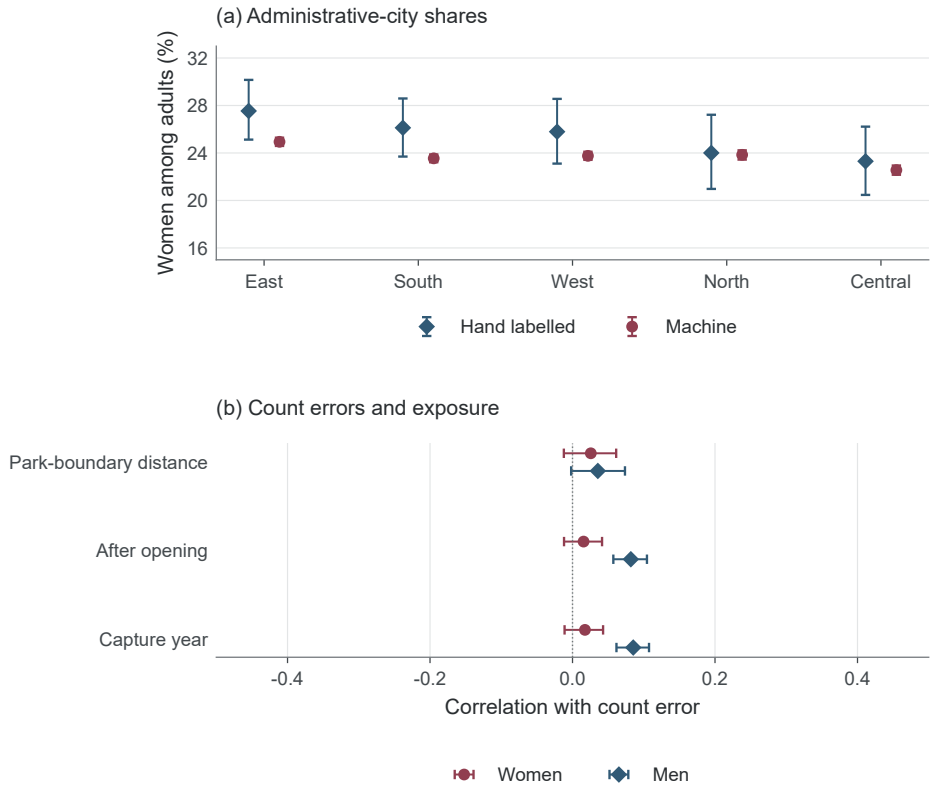
Classification errors and RPTRA exposure. For the analysis of RPTRA openings, a particular concern is whether classification errors differ systematically between streets near and far from the facilities, or change after an opening. I examine this by relating errors in predicted pedestrian counts to distance from the nearest RPTRA boundary, post-opening status and image capture year. The correlations are small for both women and men (Figure 3, panel b), suggesting that the observed classification errors are not systematically associated with the main dimensions of the empirical design.

Comparison with independent mobility data. I further validate the spatial variation in street presence using aggregate mobile-phone mobility data from Veraset for 2019.[8] Matching these data with Street View counts across 770 one-kilometre grid cells, I find a positive correlation of 0.454 between log-one-plus pedestrian intensity and log-one-plus phone-based trip-end density (arrivals plus departures). This suggests that the imagery captures meaningful spatial variation in urban activity, although the phone data do not distinguish users by gender (Appendix A.3).

Finally, I check whether the RPTRA estimates depend on how the classifier assigns gender to detected people. I compare the baseline method with an alternative approach to calibrating predicted

[8]The released aggregate flows are from the replication package for Kreindler et al. (2026); I do not access individual smartphone trajectories.

Figure 3: Spatial agreement and measurement errors



Notes: Panel (a) compares ratios of summed female and adult counts in the 2013–2019 hand-labelled and machine samples. Machine counts use sampling and observation weights; the observation frames differ. Panel (b) uses calibrated count minus hand count on 4,762 complete matched panoramas, with the reduced-form hex-centroid-to-boundary distance. Post-opening status follows the linked hand-label sample. Bars are 95 per cent location-bootstrap intervals (2,000 draws). Count-error comparisons can reflect missed or retained people; the separate crop-classification metrics condition on detection. These descriptive correlations do not test that the boundary treatment effect is unbiased.

probabilities (isotonic calibration) and with a method that assigns each person to a single category rather than using probabilities (hard labels). The estimated increases in street presence near RPTRAs remain positive for both women and men across these methods (Appendix A.4.3). Using hard labels produces more cells with zero predicted counts, particularly for women. These checks vary how detected people are classified, holding the initial pedestrian detections fixed.

Table 1: Street-presence validation

Measure	Interpretation	Women (%)	Men (%)
Precision	Predicted labels that match the audit	79.7	85.1
Recall	Audited class members correctly labelled	76.0	90.4
F1	Harmonic mean of precision and recall	77.8	87.7

Notes: The aggregate validation shares are reported for the machine classified sample. Precision, recall and F1 evaluate class assignment among hand-audited detections.

3.6 Coverage and Interpretation

The gender gap persists across the daylight periods identified from shadows. Women account for 26.3 per cent of identified adults in the morning, 24.6 per cent around midday and 23.9 per cent in the afternoon or evening. Women are also less likely than men to appear in a panorama in each period. These

comparisons cover the 61 per cent of hand-labelled panoramas for which a capture period can be assigned (Appendix A.5).

Two further measurement issues matter for the RPTRA opening estimates. First, the accuracy of pedestrian detection may vary with crowding and visibility. If an opening changes the number of people standing together or the proportion of children present, it may also change the proportion of people detected by the model. Second, Street View images can capture people beyond the street itself. Although each observation is assigned to a cell based on the camera's location, some people photographed near an RPTRA boundary may be inside the facility rather than on the surrounding street. The opening estimates therefore capture changes in activity visible from nearby streets, potentially including activity within the facilities.

4 Work and Street Presence

4.1 The Gap and Its Geography

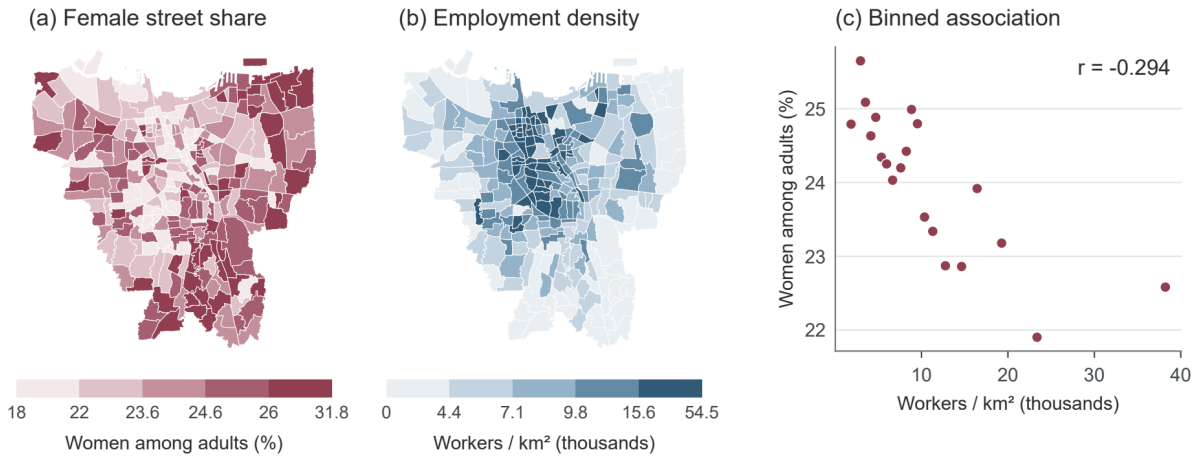
The gender gap in street presence varies considerably across Jakarta. Across the 261 mainland neighbourhoods, women account for between 18.1 and 31.8 per cent of identified adults observed on the street, with relatively high and low shares clustered geographically (Table A4). Women are particularly under-represented in central and north-western Jakarta, while their share is higher in parts of eastern and southern Jakarta. The same broad pattern appears in both the hand-labelled and machine-classified samples, with East Jakarta having the highest female share and Central Jakarta the lowest among the five administrative cities (Appendix A.4.3).

This geography is closely related to where economic activity takes place. Figure 4 compares female street presence with employment density from the 2016 Economic Census. Jobs are concentrated in central Jakarta, where women's share of street presence tends to be lower, with an unweighted neighbourhood-level Pearson correlation of -0.294. By contrast, neither residential population density nor the female share of residents closely reproduces this pattern (Appendix Figure A4). Street presence therefore reflects more than the composition of nearby residents, with the location of employment and other activities also shaping who is observed in different parts of the city.

Work is a natural starting point for understanding this spatial pattern. Men are more likely to be employed, and some occupations involve spending much of the working day in public view. Yet the gender difference in employment is considerably smaller than that in street presence. In 2018, the adult female-to-male population ratio was 1.0112, while employment rates were 45.3 per cent for women and 72.7 per cent for men: women's rate was 62.3 per cent of men's. The hand-labelled imagery records 6,029 women and 18,266 men on foot or seated, a female-to-male ratio of 0.3301, or approximately one woman for every three men.[9] The larger street gap may reflect differences in where and how long women and men work, as well as how they commute. I examine these margins through an accounting of visible working time.

[9] Sources: survey-weighted SAKERNAS 2018, mainland adults aged 15 and over; 10,216 hand-labelled panoramas, 2013–2024 excluding 2020–2021. Street counts exclude children, unidentified adults and non-motorised riders. The accounting uses 2015–2019 imagery and adjusts for the population gender ratio.

Figure 4: Female street presence and the geography of employment



Notes: Panels (a) and (b) show female street share and employment density across Jakarta's 261 mainland neighbourhoods. Female shares are calculated from sampling- and coverage-weighted machine-classified adult counts in 2013–2019. Employment density is measured using the 2016 Economic Census, in workers per square kilometre. Panel (c) plots unweighted neighbourhood means within 20 bins of 13–14 neighbourhoods, sorted by employment density. The displayed Pearson correlation, $r = -0.294$, uses all 261 unbinned neighbourhood observations with equal weights. It does not adjust for other neighbourhood characteristics.

4.2 An Accounting Framework

Employment contributes to street presence through time spent working in public view and through journeys to and from work. The amount of visible time generated by employment therefore depends not only on whether people work, but also on their working hours, workplace locations and commuting. Building on the visible-person-time interpretation in Section 3, I combine these four margins in an accounting of the gender gap:

$$V_g = \underbrace{p_g\,(h_g\phi_g + c_g)}_{\text{measured work component } W_g} + \underbrace{\ell_g}_{\text{outside measured work}}, \tag{1}$$

where V_g is visible person-time per adult of gender g, p_g is the employment-to-population ratio, h_g is working time per employed person, ϕ_g is the fraction of work visible from the street, and c_g is visible commuting time per employed person. ℓ_g, includes non-work travel and other activity in public view.

The visibility of work varies considerably across occupations and workplaces. To measure it, I use the 2023 National Labour Force Survey (SAKERNAS), which records respondents' usual work location and occupation. Work locations distinguish employment at home, at another person's home, in offices and factories, in malls and covered markets, and on streets, roadsides or in itinerant locations. I combine these locations with occupation to distinguish jobs that involve working in public view from those performed primarily indoors or inside vehicles. For example, street vendors and roadside workers are likely to remain visible for much of their working day, while workers in offices or factories contribute little visible working time. I also distinguish machine and vehicle operators from other street workers, since time spent operating a vehicle is generally not visible in the imagery.

I assign a visible fraction a^k to each work-location and occupation category k, representing the share of working time expected to be visible from the street. If $\theta^k_{g,h}$ denotes the share of total usual working

hours supplied by employed adults of gender g in category k, the hours-weighted visibility of work is

$$\phi_g = \sum_k \theta^k_{g,h} a^k.$$

Because the survey does not record the fraction of each working day spent in public view, I consider low, middle and high visibility assignments. The baseline uses the middle assignment and applies the same visible fraction to women and men within each category. Gender differences in ϕ_g therefore arise from differences in the types and locations of jobs held by women and men, rather than from assumed differences in visibility within the same category. Appendix C.4 reports the detailed crosswalk and alternative assignments.

For each employed respondent, I multiply usual working hours by the assigned visibility fraction and average across individuals using survey weights before converting to daily minutes. The resulting visibility fraction ϕ_g therefore accounts for differences in working hours as well as the distribution of workers across occupations. For 2018, I assign each respondent the average visibility fraction of their gender and occupation group estimated from the 2023 data, rather than observing visibility directly.

I compare the gender gap in observed street presence with the gap predicted by measured work-related activities. Let S_f denote the female share of observed adults and N_g the adult population of gender g. The ratio $(1 - S_f)/S_f$ gives male-to-female street presence in the imagery. To express this ratio in visible person-time per adult, I adjust for the relative population of women and men:

$$R \equiv \frac{V_m}{V_f} = \frac{N_f}{N_m} \frac{1 - S_f}{S_f}.$$

The adult population is approximately balanced by gender ($N_f/N_m = 1.011$). For the 2015–2019 imagery sample,[10] the resulting male-to-female ratio is thus $R = 2.807$: men generate approximately 2.8 times as much visible street presence per adult as women. The corresponding ratio implied by employment, hours, workplace visibility and commuting is lower, at $\rho_W = W_m/W_f = 2.261$.

To determine how much of the observed gap measured work can explain, I consider the largest possible contribution of work given these ratios. Since visible time outside measured work cannot be negative, this maximum occurs when women's visible time outside work is set to zero ($\ell_f = 0$). Any positive female exposure outside work would require greater male exposure outside work to reproduce the observed ratio. The minimum share of the gap outside measured work is therefore

$$\frac{\ell_m - \ell_f}{V_m - V_f} \geq \frac{R - \rho_W}{R - 1} \simeq 0.302.$$

Under the baseline visibility assignments, measured work thus accounts for at most 69.8 per cent of the observed gender gap, leaving a conditional minimum residual of 30.2 per cent. This minimum is obtained by setting women's visible time outside measured work to zero, holding the work-related inputs fixed. Its magnitude depends on the visibility assigned to different jobs and the treatment of unidentified adults. Appendix C provides the derivation, Appendix Table C3 reports the inputs, and

[10] I restrict the accounting to 2015–2019 to align the street-presence measure more closely with the 2018 labour-force survey and the period of RPTRA construction. Including earlier years would extend the comparison further from the labour-market reference period, while the full imagery sample also includes observations from 2022–2024, after the RPTRA rollout and the COVID-19 pandemic. The restricted window provides a more comparable reference period for the accounting, although the sources do not cover identical years.

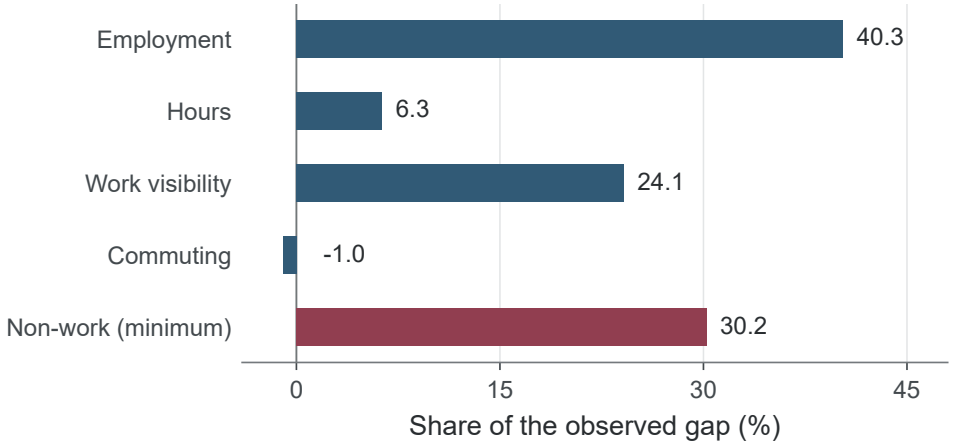
Appendix Figure C1 examines alternative visibility assignments.

4.3 Where Women Work

Gender differences in employment and workplace location are substantial. In 2018, 72.7 per cent of men and 45.3 per cent of women were employed. Among employed adults, women were more than twice as likely to work at home: 27.5 per cent compared with 13.3 per cent of men. They were also less than half as likely to work on streets, roadsides or in itinerant locations: 6.4 per cent compared with 16.3 per cent of men, according to the 2023 work-location data. Differences in working hours are much smaller, with employed men reporting 48.6 usual hours per week compared with 45.0 for women (Table A5).

I use these differences to decompose the observed street-presence gap into the four work-related margins following the methodology in Kleven and Landais (2017). Since employment, hours and workplace visibility enter multiplicatively, I allocate their interactions symmetrically across the components (Appendix C.3). Under the baseline assumption that women spend no visible time outside measured work ($\ell_f = 0$), employment accounts for 40.3 per cent of the gap, workplace visibility for another 24.1 per cent and working hours for 6.3 per cent (Figure 5). Workplace visibility thus contributes more than half as much as employment itself, reflecting the large differences in where women and men work.

Figure 5: Work-related contributions to the gender gap in street presence



Notes: Bars report shares of the male-minus-female gap in visible person-time per adult, including the negative commuting contribution and the conditional minimum residual. The symmetric decomposition averages the two gender orderings of employment and exposure before separating hours from workplace visibility. Employment and hours use survey-weighted SAKERNAS 2018 data for mainland Jakarta adults aged 15 and over; work-location shares and commuting inputs come from 2023 surveys. Street presence is measured using hand-labelled pedestrians in 2015–2019. The baseline uses the middle visibility assignment, common to women and men within each work category, and counts walk-only commuting. Unidentified adults are excluded, and the residual is evaluated at $\ell_f = 0$.

Commuting contributes in the opposite direction, offsetting the gap by approximately 1.0 per cent. In the 2023 commuter survey, 25.7 per cent of female work commuters travel entirely on foot, compared with 12.8 per cent of men (Table A5). Combining these shares with commuting durations gives slightly greater visible commuting exposure for women per employed person. The baseline calculation includes only walk-only journeys, excluding walking and waiting during mixed-mode commutes. Given women's greater reliance on walking and public modes of travel, this convention provides a conservative benchmark for the contribution of commuting to the measured gender gap. The commuter survey covers journeys

crossing administrative boundaries rather than all employed residents, and Appendix C.4 describes the construction and alternative assumptions.

4.4 Accounting Counterfactuals

How much of the difference in visible work would remain if women and men had the same employment rate? Holding hours, workplace visibility and commuting per worker fixed, equalising employment reduces the male-to-female work-exposure ratio from 2.261 to 1.41. Writing $\bar{v}_g = h_g\phi_g + c_g$,

$$\rho_W^{\text{eq}} = \frac{p_m \bar{v}_m}{p_m \bar{v}_f} = \rho_W \frac{p_f}{p_m} = 1.41. \tag{2}$$

Even with equal employment rates, men would generate 41 per cent more visible work exposure per adult than women, reflecting differences in hours, workplace visibility and commuting. This calculation holds exposure per worker fixed and does not predict how total street presence would respond to an increase in women's employment (Appendix C.2).

The conditional minimum residual also depends on how much visible time women spend outside measured work. The baseline sets this time to zero. If instead women's visible time outside work equals half their measured work exposure, the residual increases from 30.2 to 53.5 per cent of the observed gap. Appendix Table C1 reports the corresponding scenarios.

4.5 Sensitivity and Interpretation

The conditional minimum residual is sensitive to the visibility assigned to different jobs. Reducing women's assigned visible working time by 21.3 per cent, while holding male inputs and commuting fixed, eliminates the positive residual under the baseline calculation (Figure C2). Alternative assignments of unidentified adults also affect the result (Section C.6). Repeating the exercise with SAKERNAS 2023 and the 2022–2024 imagery gives a conditional minimum residual of approximately 39 per cent (Appendix C.7).

The accounting combines work-related exposure measured for Jakarta residents with street presence observed among everyone visible in the city, including visitors and non-resident workers. Differences in population coverage and unmeasured work exposure may therefore contribute to the residual alongside activity outside work. Moreover, care responsibilities, safety concerns and social norms may affect several components of visible time simultaneously. I next examine gender differences in overall trip-making and in the destinations of journeys outside work using household travel diaries.

5 Travel and Street Presence

Having examined how market work contributes to gender differences in street presence, I turn to travel as another dimension of how women and men use the city. Using the 2002 household travel diaries, I first compare overall trip-making with street presence. While the imagery captures who occupies public space, the travel diaries record how frequently people travel, for what purposes and to which destinations. I then focus on activities outside work, examining gender differences in the frequency and spatial reach of non-work journeys. These patterns are particularly relevant to understanding how the provision of nearby destinations may shape women's and men's use of the city.

5.1 Total Trip-Making

Women make fewer trips than men overall, but the difference is much smaller than the gap in street presence. Across all purposes, women record 82 per cent as many trip legs per adult as men, compared with a female-to-male ratio of approximately 33 per cent in the street imagery (Table 2). Women are also less likely to make any trip during the day and spend less time travelling overall. Although these data cover different periods, the contrast between trip frequency and street presence is striking: differences in how often women and men travel are much smaller than differences in how frequently they are observed in public view.

Table 2: Trip-making, trip frequency and journey travel time

	Trip-making (%)		Legs per adult	
	Men	Women	Men	Women
Any trip	85.8	74.4	2.366	1.938
Work and business	59.2	23.7	0.803	0.283
School	14.1	13.5	0.147	0.141
Shopping	8.1	31.2	0.085	0.327
Private	23.8	20.8	0.295	0.248
Other	6.9	6.4	0.082	0.074
Non-work	35.5	48.6	0.463	0.648

Completed journey	Travel minutes		Adjusted p
	Men	Women	
Work	90.2	90.9	0.362
School	82.2	83.4	0.328
Shopping	56.1	38.3	< 0.001
Business	104.5	81.3	< 0.001
Private	57.8	54.2	< 0.001
Other	47.6	44.7	0.038

Notes: JICA 2002 survey data. The top panel reports survey-weighted trip participation among 185,744 DKI Jakarta residents aged 15 and above, including non-travellers. Non-work trips include shopping, private and other purposes, excluding return-home legs. Any trip includes return-home legs. The bottom panel reports travel times for completed home-to-home, single-purpose journeys, including return travel but excluding time spent at destinations. Person weights are used throughout, with standard errors clustered by residential survey zone.

The two measures capture different aspects of urban activity. Travel diaries describe movement between places, while street imagery also records time spent working, waiting and undertaking other activities in public spaces. Together, they allow me to examine both the journeys women and men make and where they are observed across the city.

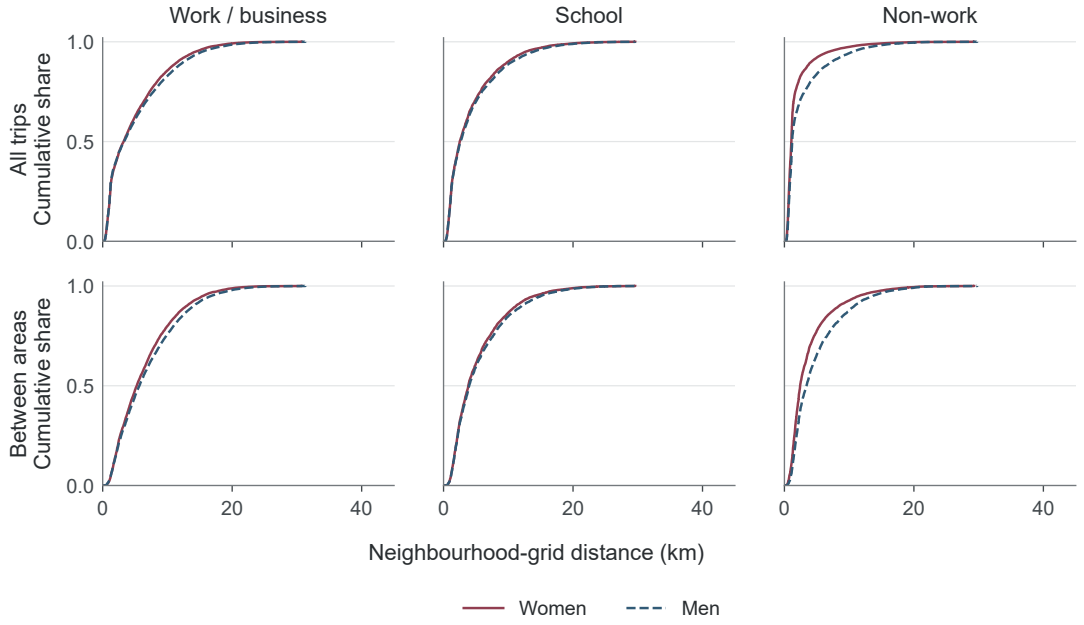
5.2 Non-Work Trips and Spatial Reach

The gender difference reverses for travel outside work. Women make approximately 40 per cent more shopping, private and other non-work trip legs per adult than men, despite making fewer trips overall. Nearly half of women undertake at least one such trip during the survey day, compared with about one-third of men. Among those making non-work trips, however, the number of legs is similar across genders, at around 1.3 per person. The higher frequency of women's non-work travel therefore comes primarily from a greater likelihood of undertaking these activities, rather than more journeys among those who participate.

These differences in travel purpose are accompanied by differences in destination choice. Women's

non-work activities are more concentrated near home, consistent with evidence that household production-related travel tends to be more local (Sánchez de Madariaga, 2013; Sánchez de Madariaga and Zucchini, 2019; Ng and Acker, 2018; Gauvin et al., 2020). Figure 6 compares the distance distributions of trips originating at home. Women's non-work journeys are shorter across much of the distribution, including among trips that cross neighbourhood boundaries. For these journeys, women travel an average of 3.79 kilometres compared with 4.86 kilometres for men. In contrast, the distributions of work and school trips are much closer. Women's greater concentration in nearby destinations is therefore most pronounced for activities outside work.

Figure 6: Trip-distance distributions by gender and purpose



Notes: Survey-weighted cumulative distributions of trip distances from the 2002 SITRAMP survey, covering 167,396 trips departing from home by residents aged 15 and above. Trips are grouped into work/business, school and non-work purposes, with all transport modes included. The upper row includes trips within and between neighbourhoods; the lower row includes only trips between neighbourhoods. Distances are calculated between neighbourhood centroids, with within-neighbourhood distances approximated using an equal-area circle. Return-home trips and destinations outside the 261 mainland neighbourhoods are excluded. Distances are geographic approximations rather than actual travel distances.

Travel time provides a similar picture. A completed shopping journey, including the return home, takes women an average of 38 minutes of travel compared with 56 minutes for men. By contrast, work and school journeys take approximately the same amount of time across genders (Appendix Table B1). These differences suggest that the greater localness of women's non-work activity is not confined to the way distances are measured: it is also apparent in the time spent travelling.

The pattern is not uniform across all non-work purposes. Recreation alone shows a different gender ordering in average distance from the broader non-work category. Appendix Tables E1 and B2 report the purpose-specific distances and trip-chaining comparisons, while Appendix Figures B1 and B2 map the origin–destination flows.

Taken together, the travel patterns show that women make fewer journeys overall but undertake more non-work trips, with destinations that tend to be closer to home. The greater concentration of women's non-work activity near home raises the question of how the availability of local destinations

affects their use of the city. I examine this through Jakarta's expansion of neighbourhood public spaces, first estimating changes in street presence around newly opened facilities and then modelling how their locations affect trip-making.

6 Local Public Spaces

The travel evidence shows that women's non-work journeys are more concentrated near home, making the provision of local destinations particularly relevant. I examine the effect of new neighbourhood public spaces using Jakarta's RPTRA programme, described in Section 2.2. I first estimate how openings change the number of women and men visible near the facilities. I then compare the gender composition of people photographed inside RPTRAs with those observed on surrounding streets, drawing on the hand labels to examine differences in street activities.

6.1 Design and Identification

Of the registered RPTRAs, about 86 per cent opened between 2015 and the end of 2017. Their locations and opening dates depended partly on the availability of public land and implementation decisions rather than random assignment. Comparing neighbourhoods that received an RPTRA with those that did not could therefore confound the effects of the programme with differences in the characteristics and development of those neighbourhoods. Instead, I compare changes in street presence close to each RPTRA with changes on streets farther from the same facility, before and after it opens. This approach uses the staggered timing of openings and spatial variation in exposure within each facility's surrounding area.

The unit of observation is a hexagonal cell 50 metres wide. In the baseline specification, I assign each cell to its nearest RPTRA, measured from the cell centroid to the facility boundary, and retain this assignment throughout the study period. Each cell therefore has a single reference facility, even when other RPTRAs are nearby or open at different times. The assigned facility's opening date determines treatment timing, while distance from its boundary determines the cell's exposure band.

I measure distance from a cell's centroid to the nearest point on an RPTRA boundary, setting it to zero if the centroid lies inside the facility. In the baseline specification, d_i denotes this distance for the nearest RPTRA. Mapped boundaries are available for 298 sites; for the remaining 17, I approximate the boundary using a circle based on the registered area and location.[11] I divide cells into six distance bands within 750 metres of the boundary: 0–25, 25–50, 50–150, 150–250, 250–500 and 500–750 metres. Cells between 750 and 1,500 metres form the comparison group. The multiple bands allow the response to vary with proximity rather than imposing a common effect on all streets within a single treatment radius. Since the cells themselves are 50 metres wide, the innermost band captures the immediate boundary environment rather than a precisely measured 25-metre radius. Approximately 94 per cent of observations in this band come from cells that touch or overlap the facility boundary.

In an alternative specification, each cell can enter the catchment of multiple RPTRAs within 1,500 metres. Its distance band and treatment status are then defined separately for each facility, allowing streets to contribute to more than one RPTRA-specific comparison.

[11] The mapped sites have a median area of approximately 1,600 square metres, equivalent to a circle with a radius of 22.5 metres. Reported locations lie a median of 13.9 metres from polygon centres.

This is a spatial difference-in-differences design, in which nearby streets serve as comparisons for changes around the same facility. Butts (2023) examines difference-in-differences designs with spatial spillovers, highlighting the importance of comparable untreated trends across distance bands and the possibility that comparison areas are themselves affected by treatment. The latter is particularly relevant for RPTRAs, since an opening may attract pedestrians to surrounding streets beyond the facility's immediate boundary. I therefore estimate changes across several distance bands and examine how the results vary with the choice of comparison ring. As discussed in Butts (2023) and Clarke (2017), effects extending into the comparison area would make the estimates relative responses rather than the full effect of an opening.

For each cell i, I use images captured in year t and month m. The primary sample contains observations from 2013, 2015, 2018 and 2019, with 184,410 cell-year observations from 62,269 cells assigned to 315 facilities. I exclude 2014 because imagery is available for very few cells and none changes treatment status in that year. I also exclude 2016 and 2017, when most facilities opened, because imagery availability differs systematically across distance bands and is correlated with observed street activity.[12] After excluding separated observations and singleton fixed-effect groups, the Poisson estimation sample contains 118,588 observations around 314 RPTRAs (Table 3).

I estimate the effect of RPTRA openings on pedestrian counts using Poisson pseudo-maximum likelihood (PPML) (Santos Silva and Tenreyro, 2006). I specify the conditional mean of pedestrian counts as an exponential function of the explanatory variables, ensuring that fitted conditional means are positive and allowing the effects of openings to be interpreted as proportional changes in pedestrian presence. The specification also accommodates cells with zero pedestrians and the fractional expected counts generated by the classifier:

$$\mathbb{E}[y_{itm}] = \exp\left(\alpha_i + \delta_{p(i)t} + \xi_m + \sum_{b=1}^{6} \beta_b \, \mathbf{1}[d_i \in b] D_{itm}\right). \tag{3}$$

The dependent variable y_{itm} is the expected count of female or male pedestrians in cell i, captured in year t and month m. The treatment variable is the interaction between D_{itm}, an indicator equal to one once the cell's assigned RPTRA has opened, and $\mathbf{1}[d_i \in b]$, an indicator for the cell's distance band. The coefficients β_b measure how pedestrian presence in each distance band changes after opening relative to the change in the 750–1,500 metre comparison ring around the same RPTRA. I estimate all six bands jointly and report the effects as percentage changes, $100[\exp(\beta_b) - 1]$. Cell fixed effects α_i absorb permanent differences between streets, RPTRA-by-year fixed effects $\delta_{p(i)t}$ absorb changes shared by streets around the same facility in a given year, and capture-month fixed effects ξ_m account for seasonal variation. The regressions use sampling and imagery-coverage weights, with standard errors clustered at the RPTRA level.

For the alternative specification, I retain the same distance bands but allow cells to enter multiple facility-specific comparisons. To prevent cells associated with multiple RPTRAs from receiving greater total weight because they appear more than once, I divide their sampling and imagery-coverage weights by the number of facility memberships. I replace cell fixed effects with cell-by-RPTRA fixed effects and

[12] The 2014 imagery covers only 929 cells. During 2016 and 2017, the distribution of available imagery changes substantially between streets close to RPTRAs and those in the comparison ring. Including these years would make it difficult to separate changes in pedestrian presence from changes in which streets are observed. Observations with unknown capture month are also excluded.

include RPTRA-by-year and capture-month fixed effects, clustering standard errors by both physical cell and RPTRA. The alternative specification therefore changes which streets contribute to each facility's comparison, but does not eliminate potential exposure to other RPTRA openings.

Identification comes from changes in the spatial gradient of street presence around each facility. By including RPTRA-by-year fixed effects, I compare streets at different distances from the same RPTRA within a given year, rather than relying on comparisons between neighbourhoods with different opening dates. This absorbs shocks common to a facility's surrounding area, including broader changes in local economic activity or pedestrian traffic. The identifying assumption is that, without an RPTRA opening, pedestrian presence in the inner bands and the comparison ring would have changed in the same proportion.

Assigning each cell to its nearest RPTRA does not rule out exposure to other facilities. A cell may be close to several RPTRAs that opened at different dates, so its pre-opening observations relative to the assigned site may already reflect activity generated by another facility. Likewise, an opening may redirect journeys from an existing RPTRA rather than generate new activity. I examine overlapping exposure by controlling for other open sites and distance to the second-nearest facility, and by varying the sample to restrict exposure to other openings (Appendix D.6). These checks address the sensitivity of the estimates to nearby facilities, although they do not fully identify the contribution of each opening when catchment areas overlap.

Two features of the setting warrant particular attention. First, the programme was rolled out over several years, and treatment effects may differ across opening cohorts or evolve over time. Recent difference-in-differences research highlights how heterogeneous effects and treatment timing can complicate the interpretation of conventional fixed-effects estimates (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021). My preferred specification includes facility-by-year effects and estimates a distance-specific change relative to the surrounding comparison ring, rather than treating earlier-opened facilities as controls for later-opened ones. I also examine a staggered difference-in-differences specification and cohort-specific event-study estimates as supplementary checks.

To assess sensitivity to the choice of comparison ring, I repeat the estimates using alternative distances, restrict comparisons to common sets of facilities, and examine responses in adjacent bands. These exercises help assess the sensitivity of the estimated spatial profile, although they cannot by themselves distinguish spillovers from differences in untreated trends. The event-study and comparison-ring results are presented in Section 6.2, with additional specifications in Appendices D and D.6.

Finally, streets close to RPTRAs differ from those farther away in several baseline characteristics (Appendix Table D1). The cell fixed effects account for these differences in levels, while the RPTRA-by-year effects absorb shocks shared within a facility's catchment.

6.2 Street Presence at the Boundary

RPTRA openings are associated with a substantial increase in pedestrian presence on nearby streets. Within 25 metres of the facility boundary, expected female counts rise by 31.3 per cent and male counts by 23.1 per cent, relative to streets 750–1,500 metres from the same RPTRA (Table 3 and Figure 7). Applying these percentage changes to the pre-opening means of 0.266 women and 0.938 men per cell gives approximately 0.08 additional women and 0.22 additional men per observed cell. The estimates measure changes in the number of people visible at a given moment, rather than the number of daily

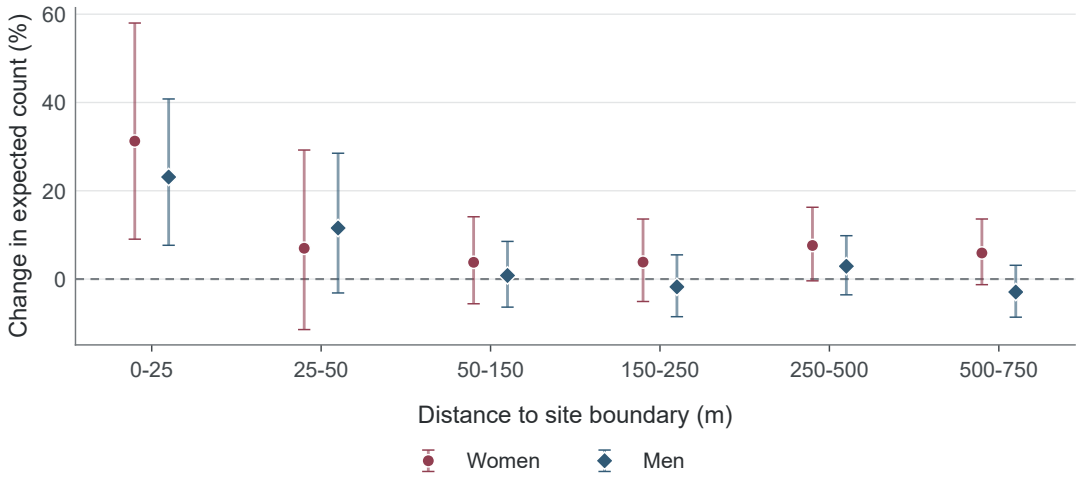
visitors.

Table 3: RPTRA openings and street presence, by distance from the boundary

	Baseline: nearest RPTRA		Multiple RPTRAs	
Boundary distance (m)	Women	Men	Women	Men
	(1)	(2)	(3)	(4)
0–25	31.3***	23.1***	32.9**	28.3***
	[12.4]	[8.4]	[15.3]	[9.8]
25–50	7.0	11.6	11.5	17.1*
	[10.3]	[8.0]	[13.2]	[10.3]
50–150	3.8	0.8	1.9	2.7
	[5.0]	[3.8]	[5.4]	[3.5]
150–250	3.8	-1.8	-2.6	-1.6
	[4.8]	[3.6]	[4.6]	[2.9]
250–500	7.6*	2.9	3.4	2.3
	[4.2]	[3.4]	[2.8]	[2.1]
500–750	5.9	-2.9	4.8*	1.4
	[3.8]	[3.0]	[2.7]	[2.0]
Fixed effects	Cell	Cell	Cell–park	Cell–park
RPTRA × year FE	Yes	Yes	Yes	Yes
Capture-month FE	Yes	Yes	Yes	Yes
Clustering	RPTRA	RPTRA	RPTRA + cell	RPTRA + cell
Estimation rows	118,588	118,588	524,776	524,776
Physical cells	34,186	34,186	34,185	34,185
RPTRAs	314	314	315	315
Pre-opening mean, 0–25 m	0.266	0.938	0.277	0.994

Notes: Entries report percentage changes in expected pedestrian counts from PPML regressions, $100[\exp(\widehat{\beta}) - 1]$, with delta-method standard errors in brackets. Columns (1)–(2) use the baseline nearest-RPTRA assignment, while columns (3)–(4) allow cells to enter the catchments of multiple RPTRAs within 1,500 metres. All specifications include six distance bands, with 750–1,500 metres as the comparison group, and use sampling and imagery-coverage weights. Columns (1)–(2) control for cell, RPTRA-by-year and capture-month fixed effects, with standard errors clustered by RPTRA. Columns (3)–(4) control for cell-by-RPTRA fixed effects, split weights across park memberships and cluster by cell and RPTRA. The sample covers 2013, 2015, 2018 and 2019. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 7: RPTRA openings and street presence, by distance from the boundary



Notes: Points report $100[\exp(\widehat{\beta}_b) - 1]$ from the gender-specific weighted PPML regressions in Equation (3), with 95 per cent confidence intervals. Distance is measured from the nearest RPTRA boundary. All six treatment bands enter jointly, with the 750–1,500 metre ring omitted. Standard errors are clustered by RPTRA.

The response is concentrated immediately around the facility. In the 25–50 metre band, estimated increases fall to 7.0 per cent for women and 11.6 per cent for men. Farther away, the estimates are smaller and generally imprecise, with coefficients beyond 50 metres jointly indistinguishable from zero. Although the point estimates suggest a sharp decline in the response with distance, the difference between the first two bands is only marginally statistically significant for women and not significant for men. The evidence therefore points to a highly localised increase in street presence.

The larger point estimate for women does not translate into a statistically distinguishable gender difference. In a stacked specification, the male-minus-female difference in the opening coefficient within 25 metres is -0.064 log points, with a 95 per cent confidence interval of $[-0.205, 0.077]$ (Table D2 and Figure D11). The estimates are consistent with an increase in women's absolute street presence, but provide little evidence that RPTRA openings change the gender composition of pedestrians on surrounding streets.

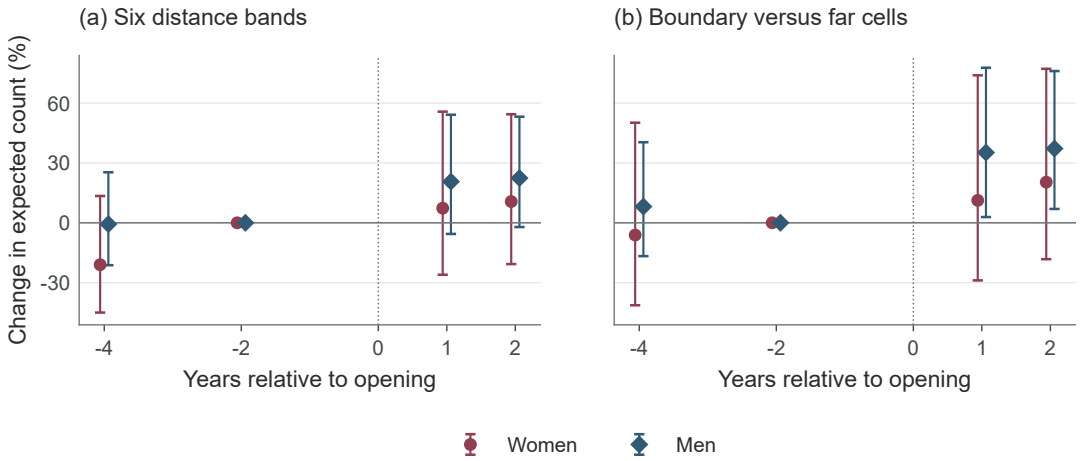
Timing and spatial comparisons. The identifying assumption requires that street presence near and farther from each RPTRA would have followed similar proportional trends in the absence of an opening. I examine this using event studies for the 2017 opening cohort, with observations from 2013, 2015, 2018 and 2019 and 2015 as the reference year (Figure 8). Panel (a) allows the response to vary across the six distance bands in the baseline specification. Within 25 metres, estimated increases range from 7.3 to 10.7 per cent for women and from 20.6 to 22.4 per cent for men across the two post-opening years. None is individually statistically significant, and a joint test does not reject equality of the gender responses ($p = 0.635$). Panel (b) restricts the sample to the 0–25 metre band and the 750–1,500 metre comparison ring. Here, the male estimates are larger and statistically significant, but the gender difference remains statistically indistinguishable from zero (joint $p = 0.445$). With only one pre-opening contrast, the event studies provide limited evidence for assessing parallel trends.

Appendix Table D9 reconciles these estimates with the pooled results. Restricting the six-band baseline to the 2017 cohort yields increases of 22.9 per cent for women and 21.9 per cent for men. The relative magnitudes change further when the sample is restricted to the boundary and comparison rings and when 2015 alone defines the reference period. The female and male estimation samples coincide within each specification, and direct tests do not reject equality of their responses.

The results also depend on the choice of comparison streets. I re-estimate the boundary response using alternative distance rings while holding the treated band fixed at 0–25 metres. The estimated effects remain positive, but are larger when using streets 1,000–1,500 metres away than when using those 750–1,000 metres away, even after restricting the sample to the same facilities (Figures D7 to D9). This variation may reflect spillovers into the comparison areas or differences in how streets at different distances would have evolved without an opening. The comparison-ring sensitivity therefore affects the magnitude of the estimated response, although not its sign (Butts, 2023; Clarke, 2017).

An increase in pedestrian counts immediately around a facility could also reflect activity redirected from nearby streets. However, I find no corresponding decline in the adjacent 25–50 metre band. This provides no clear evidence of displacement from those streets, although activity may have been redirected from elsewhere. I also compare pedestrian counts with motorcycle riders, whose presence increases significantly less following RPTRA openings. The rider-minus-pedestrian difference is -0.134 log points ($p = 0.041$). Although motorcycle traffic may itself respond to an opening, the smaller increase among

Figure 8: Street presence around RPTRA openings: the 2017 cohort



Notes: Points report percentage changes in expected pedestrian counts within 25 metres of RPTRAs opened in 2017, relative to the 750–1,500 metre comparison ring, with pointwise 95 per cent confidence intervals. Event year -2 (2015) is the reference. Panel (a) includes all six distance bands, with separate event-time interactions for each, using 75,987 observations per gender around 202 RPTRAs. Panel (b) includes only the 0–25 metre band and the comparison ring, using 20,722 observations around 172 RPTRAs. Both specifications include cell, RPTRA-by-year and capture-month fixed effects, use sampling and imagery-coverage weights, and cluster standard errors by RPTRA. The sample covers 2013, 2015, 2018 and 2019.

riders suggests that the pedestrian response is not merely part of a general increase in road activity.

Additional exposure and measurement checks. To assess exposure to nearby RPTRAs, I control for other open facilities and distance to the second-nearest site. The estimated boundary effects remain broadly similar (Table D5). Since comparison cells are at least 750 metres from every RPTRA under the nearest-site assignment, direct exposure to other facilities is less likely in the comparison ring, although effects extending beyond this distance could still influence the estimates.

I also examine the robustness of the results to alternative estimation samples and outcome definitions. The estimated boundary responses remain positive when restricting the analysis to cells observed in multiple periods, excluding particular cohorts or facilities, using alternative estimators and replacing predicted counts with indicators for whether any pedestrians are present (Sections D and D.6). The results are also similar under alternative methods of assigning gender to detected pedestrians. District-level score-bootstrap tests support positive boundary responses for both genders but do not reject equality between their proportional effects.

Measurement error remains a concern because the outcomes rely on automated pedestrian detection. As discussed in Section 3.5, classification errors in the audited sample are weakly correlated with distance to the RPTRA boundary, opening status and capture year (Figure 3). However, if openings change crowding or visibility, pedestrian detection rates may also change, which the gender-classification audit cannot identify. Moreover, Street View records people visible from the camera rather than exclusively those standing on the road. Some people photographed near an RPTRA boundary may therefore be inside the facility itself. The estimated response concerns activity visible from nearby streets, potentially including facility users in view.

An increase in street presence could also reflect changes in the residential population around RPTRAs rather than changes in how existing residents use nearby streets. I examine registered population and assessed land values to assess this possibility (Appendices D.8 and D.9). Population increases emerge

mainly after 2020, several years after most RPTRAs opened and outside the main street-imagery period. These estimates are sensitive to specification and changes in administrative population records. I also find no consistent evidence of increases in assessed land values near the facilities, although these assessments are imperfect measures of market prices. Together, these results provide limited support for residential sorting as the primary explanation for the earlier increase in street presence, without ruling out changes in residential location or capitalisation over a longer period.

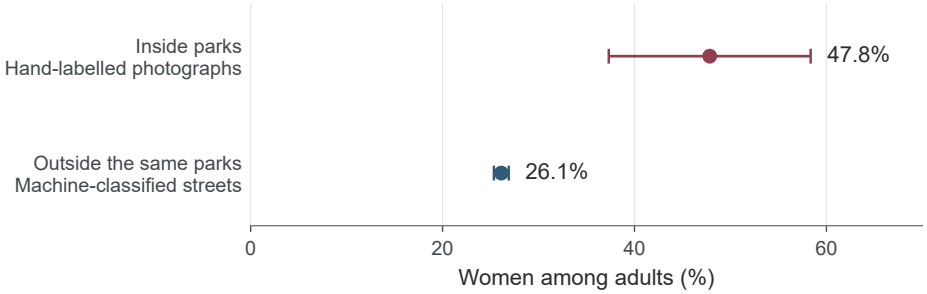
The boundary estimates therefore capture an increase in visible activity near RPTRAs, without distinguishing new journeys from activity redirected from elsewhere. Since RPTRAs combine recreational space with staffing and community services, the estimates capture the response to the facility as a whole rather than to any individual feature. The next subsection examines the gender composition of people photographed inside RPTRAs.

6.3 Inside the Spaces

The opening estimates show that RPTRAs increase nearby street presence for both women and men, without a statistically distinguishable difference in their proportional responses. However, the street cameras capture only part of the activity associated with these facilities, much of which takes place inside their boundaries. To examine this activity, I hand-label 343 user-contributed Google Maps photographs from 208 RPTRAs. I distinguish photographs depicting ordinary use from organised events and posed groups, which can differ substantially in the people they capture. The main analysis focuses on 217 ordinary-use photographs from 161 facilities, weighted by their audit sampling probabilities.

I compare the gender composition of adults in hand-labelled interior photographs with calibrated machine-classified street counts from post-opening observations within 250 metres of the same RPTRA boundaries. The matched sample covers 152 facilities. Women account for 47.8 per cent of identified adults inside RPTRAs and 26.1 per cent on surrounding streets, a difference of 21.7 percentage points (95 per cent confidence interval $[11.3, 32.1]$; $p < 0.001$) (Figure 9 and Table 4).

Figure 9: Female representation inside RPTRAs and on surrounding streets



Notes: Both samples cover the same 152 RPTRAs. Interior shares use 202 hand-labelled ordinary-use photographs, weighted by inverse audit sampling probabilities. Street shares use Dirichlet-calibrated expected women divided by expected women plus men, weighted for sampling and coverage, in post-opening 50-metre cells with centres outside and within 250 metres of the nearest park boundary (2015, 2018 and 2019). Children are excluded. Points are pooled weighted shares, with 95 per cent intervals clustered by park, conditional on classification and sampling weights. Photographs and street imagery differ in dates and selection.

The higher female representation inside RPTRAs suggests that these facilities may attract a different

Table 4: Female share inside facilities and on streets outside the same parks

	Inside: hand labels	Outside: machine counts
Female share (%)	47.8	26.1
95% interval	[37.3, 58.4]	[25.3, 26.9]
Photographs / street cell–years	202	13,861
Matched parks	152	152
Parks with positive adult counts	66	150
Inside minus outside (percentage points)	21.7	
95% interval for difference	[11.3, 32.1]	
Two-sided p-value	< 0.001	

Notes: Interior observations are ordinary-use Google Maps photographs, weighted by inverse audit-sampling probabilities. Exterior observations use calibrated expected adult counts from post-opening Street View imagery in 8,123 distinct 50-metre cells (2015, 2018 and 2019), with sampling and imagery-coverage weights. Cell centres lie outside and within 250 metres of the nearest RPTRA boundary, using mapped polygons where available and area-based circular approximations otherwise. Both samples are restricted to the same 152 RPTRAs. Female shares are calculated from pooled weighted counts of women and men, excluding children, without weighting parks equally. Observations with no detected adults contribute zero to the totals.

mix of activities and users from those observed on surrounding streets. RPTRAs combine playgrounds, community rooms and services for children and families, many of which involve spending time within the facility rather than on nearby roads. The difference in composition is therefore consistent with women's participation in these public spaces being greater than their representation on surrounding streets would suggest. Since the photographs are selected uploads, however, they do not measure representative attendance or establish how much additional activity the openings generate.

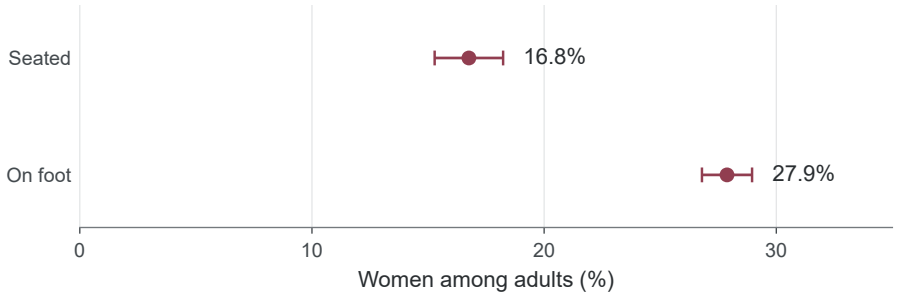
6.4 Activities on the Street

The preceding results present a contrast: women are more equally represented inside RPTRAs than on surrounding streets, yet the estimated proportional responses to openings are not statistically distinguishable by gender. One possible explanation lies in the relationship between trips and time spent in public. Street View measures the number of people present at a particular moment, which depends not only on how frequently people visit an area but also on how long they remain visible there. A facility that attracts more female visitors need not produce a larger increase in female street presence if women spend less time on surrounding streets than men.

The hand-labelled street images provide suggestive evidence of differences in how women and men use these spaces. Women account for only 16.8 per cent of seated adults, compared with 27.9 per cent of adults recorded as on foot (Figure 10). Women's under-representation is therefore particularly pronounced among those sitting in public, consistent with gender norms that may discourage women from lingering or loitering on streets. However, the on-foot category includes people standing as well as walking, and the images do not measure how long individuals remain in either activity.

Other activities may reinforce these differences. Women are more frequently observed buying goods and accompanying children or other adults, while men are more frequently observed selling or working as street vendors (Panel A of Figure A6). These patterns suggest that women's street presence may be more closely associated with shopping and accompanying others, while men's presence may reflect more activities that involve remaining in one place for extended periods.

Figure 10: Female representation among seated adults and adults on foot



Notes: Women divided by women plus men within each posture category, using 10,216 hand-labelled street panoramas at 2,963 locations in 2013–2024, excluding 2020–2021. Seated counts are 1,121 women and 5,568 men; on-foot counts are 4,908 women and 12,698 men. Bars show 95 per cent intervals from the ratio estimator, clustered by location. Counts are unweighted and exclude children, unidentified adults and non-motorised riders. On foot includes standing. These citywide snapshots measure composition, not visit duration or effects of RPTRA openings.

This distinction is particularly relevant to RPTRA openings. New facilities may attract visitors while also creating opportunities for street vendors on surrounding streets. If women primarily travel to the facility and spend much of their visit inside, while male vendors remain outside for several hours, the resulting increase in street presence could be similar across genders even if the facility attracts a more gender-balanced group of users. This could help explain why higher female representation inside RPTRAs need not translate into a larger proportional increase in female street presence. The activity patterns are citywide and descriptive, however, and do not establish whether RPTRA openings actually increased vending or changed the duration of women's and men's visits.

These patterns suggest that gender differences in street presence reflect not only how often women and men travel, but also what they do at their destinations and how long they remain visible. A neighbourhood amenity may therefore attract substantial female use without producing a disproportionately large increase in women's street presence. The opening estimates also cannot distinguish additional journeys from trips redirected from other destinations, or measure the benefits of improved access. I turn to a model of trip-making and destination choice to examine these margins and evaluate how the location of public spaces affects their use.

7 A Model of Trip-Making and Destination Choice

The reduced form analyses show that RPTRAs increase the number of people visible near their boundaries, but cannot tell whether these people make additional trips or redirect visits from existing destinations. Nor do they show where a fixed quantity of public-space land would generate the greatest increase in activity. These questions require modelling the destinations residents can reach and their decisions both to undertake an activity and to choose where to do it. The distinction matters for interpreting the gender results: women may use the facilities extensively without generating a correspondingly larger increase in street occupancy, which also reflects work and time spent outside.

I develop a two-stage model of destination choice and trip-making. Conditional on making a trip for an activity, adults choose among destinations that differ in attractiveness and travel distance. The resulting measure of access enters the decision to undertake that activity. The model draws on random-utility

choice methods (McFadden, 1974) and the spatial-choice and gravity structure used in quantitative urban economics (Ahlfeldt et al., 2015; Redding and Rossi-Hansberg, 2017; Miyauchi et al., 2025). Unlike a spatial general-equilibrium model, I hold residences, employment and existing destinations fixed. I estimate travel parameters from the 2002 diaries and use them to evaluate added facilities and alternative allocations of a fixed quantity of land. The predicted outcome is additional trips for a specified activity, not total street occupancy or welfare.

7.1 Model

The city consists of $I = 261$ neighbourhoods. Neighbourhood i contains N_i^g adults of gender $g \in \{f, m\}$. Each adult may undertake activities indexed by q, for example shopping or recreation. For each activity I distinguish two margins: whether the adult makes at least one qualifying trip, and the choice of destination j conditional on such a trip. Individuals can travel for other purposes even when they make no trip for purpose q. Individuals choose destinations based on their attractiveness and distance from home. The decision to undertake a trip depends on access to these destinations as well as individual and household characteristics. The two decisions are estimated sequentially.

Destination choice. Conditional on undertaking activity q, a resident of neighbourhood i receives utility from destination j given by

$$u_{ijq}^g = \ln A_{jq} - \delta_q^g d_{ij} + \varepsilon_{ijq}^g, \tag{4}$$

where $A_{jq} > 0$ is the destination's attractiveness for activity q, d_{ij} is distance and $\delta_q^g \geq 0$ captures the gender-specific cost of distance in utility units. The shocks ε_{ijq}^g are independent type-I extreme-value draws, normalised to unit scale. Following the standard conditional-logit (McFadden, 1974), the probability of choosing destination j is

$$\pi_{ijq}^g = \frac{A_{jq} e^{-\delta_q^g d_{ij}}}{\Phi_{iq}^g}, \qquad \Phi_{iq}^g = \sum_k A_{kq} e^{-\delta_q^g d_{ik}}. \tag{5}$$

The denominator sums the distance-discounted attractiveness of every available destination. I assume common destination attractiveness for women and men in the baseline. Gender differences in conditional destination choice thus operate through δ_q^g. This parameter captures both monetary and non-monetary travel costs and preferences for proximity; the data does not separately identify safety, travel time or other sources of distance sensitivity. Gender-specific destination effects are examined as a robustness check.

The inclusive value $\ln \Phi_{iq}^g$ summarises how easily residents of neighbourhood i can reach destinations for activity q. It depends on the number and attractiveness of available destinations and their distance from home, with distance weighted by gender-specific travel sensitivity. Under the random-utility model, it also represents the expected utility from the best available destination, up to a constant. Adding a new destination increases this measure of access, with a larger proportional improvement where existing accessible destinations are limited. I use this measure to link destination choice to the decision of whether to undertake the activity.

Trip-making. Access to more attractive or nearby destinations may increase the likelihood that an individual undertakes an activity. I model this extensive margin as the probability that adult n of gender

g, living in neighbourhood i, makes at least one trip for purpose q:

$$P_{nq}^{g} = \Lambda\Big(c_q^g + \lambda_q^g \ln \Phi_{iq}^g + X_n' \beta_q^g + \mu_{k(n)q}\Big), \tag{6}$$

where $\Lambda(z) = e^z/(1+e^z)$ is the logistic function. The dependent variable is the probability of making a qualifying trip, while the main explanatory variable is the inclusive value $\ln \Phi_{iq}^g$, which measures access to destinations for that activity. The parameter λ_q^g captures how trip-making responds to access, allowing this relationship to differ between women and men. The vector X_n includes household characteristics, and $\mu_{k(n)q}$ are sub-district fixed effects. I estimate the equation separately by gender.

The logit specification implies that a one-unit increase in log access raises the log odds of making a trip by λ_q^g, holding other characteristics fixed. The model allows an improvement in access to increase the probability of undertaking an activity without requiring the same proportional increase in trip-making for every individual.

The model determines whether an individual makes a trip for activity q, but holds the number of trips conditional on making at least one trip fixed. I assume that individuals who undertake the activity make $\bar{m}_q^g$ trips per day, equal to the observed gender-specific mean. Expected trips per adult are therefore $P_{nq}^g \bar{m}_q^g$. A new facility can increase total trips by encouraging more individuals to undertake the activity, but the model does not allow existing trip-makers to make more trips or change how long they spend at their destinations. Since activities are modelled separately, an additional trip for purpose q does not necessarily imply that the individual would otherwise have remained at home.

7.2 Estimation

The main counterfactual focuses on activities related to the services provided by RPTRAs. Since these facilities combine recreational spaces with early-childhood and health services, I construct a composite of recreation, sports and park trips, together with private trips to educational and medical destinations. The composite approximates the range of activities supported by RPTRAs, although the survey does not directly identify childcare, accompanying children or visits to RPTRAs themselves. I examine alternative definitions, including broad non-work trips, in Appendix E.[13]

I use the 2002 SITRAMP travel survey, which records daily travel for Jakarta residents, including individuals who made no trips on the survey day. The sample covers 185,744 adults aged 15 and above living in the 261 mainland neighbourhoods.[14] I use the trip records to estimate destination choice and the full sample of adults, including non-travellers, to estimate trip-making. For each activity, I consider journeys that begin at home and end at a destination within mainland Jakarta, excluding return-home trips. Individuals whose trip destinations cannot be determined are excluded when this prevents me from identifying whether they made a qualifying trip. For the RPTRA bundle, the resulting sample contains 185,703 adults (91,993 women and 93,710 men).

I use the 2002 SITRAMP travel survey, which records daily trips for 185,744 adults aged 15 and above, including those who did not travel on the survey day.[15] I use trips originating from home and ending within mainland Jakarta to estimate destination choice, excluding return-home trips. To estimate

[13] I do not use recreation alone for the main counterfactual because the small number of reported recreational trips prevents reliable estimation of trip-making probabilities for the full sample.

[14] I exclude residents of neighbourhoods outside mainland Jakarta.

[15] The sample is restricted to residents of mainland Jakarta's 261 neighbourhoods.

trip-making, I retain both travellers and non-travellers, since the outcome is whether an individual makes at least one trip for a given activity. The resulting RPTRA-composite sample contains 185,703 adults (91,993 women and 93,710 men), after excluding observations with insufficient destination information to identify trip-making. I apply the relevant survey weights and approximate travel distances using neighbourhood centroids, with a separate approximation for trips within the same neighbourhood. Appendix E.1 describes the sample restrictions and distance calculations.

I apply the survey weights to account for the population represented by each individual and trip. Travel distances are approximated using the distance between neighbourhood centroids. For trips within the same neighbourhood, I use the expected distance between two randomly located points in an equal-area circle. Appendix E.1 provides the sample restrictions and distance calculations.

Destination choice. I estimate destination attractiveness and gender-specific distance sensitivity using observed trips between neighbourhoods. The destination-choice probabilities in Equation (5) imply that trips from an origin to a destination increase with its attractiveness and decline with distance. Aggregating these choices across individuals gives the following gravity equation:

$$\mathbb{E}\left[T_{ijq}^{g} \mid \alpha_{iq}^{g}, A_{jq}, d_{ij}\right] = \exp\left(\alpha_{iq}^{g} + \ln A_{jq} - \delta_{q}^{g} d_{ij}\right), \tag{7}$$

where T_{ijq}^{g} is the number of trips from neighbourhood i to neighbourhood j by gender g for activity q. The main parameters of interest are destination attractiveness A_{jq} and gender-specific distance sensitivity δ_{q}^{g}. Origin fixed effects α_{iq}^{g} account for differences in the total number of trips originating from each neighbourhood. I estimate the equation using PPML (Santos Silva and Tenreyro, 2006), including neighbourhood pairs with no observed trips. Standard errors are clustered by origin.

I recover destination attractiveness from the estimated destination fixed effects, normalising the median positive value to one. Destinations receiving no observed trips for an activity are assigned zero attractiveness, while those excluded for other estimation reasons are treated separately (Appendix E.2). Using these estimates, I construct the baseline access measure for each origin, gender and activity:

$$\widehat{B}_{iq}^{g} = \sum_{j} \widehat{A}_{jq} e^{-\widehat{\delta}_{q}^{g} d_{ij}}.$$

This measure captures the attractiveness of existing destinations, discounted by distance, and enters the trip-making equation in the second estimation stage.

Trip-making. In the second stage, I estimate how access to destinations relates to the probability of making at least one trip for a given activity. I estimate Equation (6) separately by gender using a person-weighted logit model, with the estimated baseline access measure $\widehat{B}_{iq}^{g}$ from the first stage. The regression controls for household income, vehicle ownership, household size, the presence of children under 15 and sub-district fixed effects.

I estimate λ_{q}^{g} using variation in access across neighbourhoods within sub-districts, controlling for household characteristics and local differences through fixed effects. The counterfactual applies this estimated relationship to changes in access generated by new facilities, assuming that it captures how trip-making responds to improved accessibility. While residential sorting and unobserved neighbourhood characteristics may influence the estimated relationship, the model provides a separate assessment of

trip-making responses without calibrating λ_q^g to the RPTRA reduced form estimates.

I estimate the access-response coefficients λ_q^g separately by gender. For the counterfactuals, I restrict these coefficients to lie between zero and one. The lower bound ensures that adding destinations cannot reduce trip-making, while the upper bound ensures that additional trips do not exceed the number of visits to the new facilities, so predicted trips redirected from existing destinations remain non-negative. Both coefficients satisfy these restrictions for the main RPTRA activity composite, leaving the point estimates unchanged.

As a robustness check, I re-estimate destination attractiveness and distance sensitivity using only trips between non-adjacent neighbourhoods. This tests whether the estimated relationship between access and trip-making depends on very local travel patterns. I then reconstruct neighbourhood access using these estimates, retaining baseline attractiveness for destinations that cannot be estimated from the restricted sample (Appendix E.3).

I use 200 bootstrap replications to account for uncertainty in both stages of estimation. In each replication, I resample origin neighbourhoods within sub-districts, jointly for women and men, and re-estimate the destination-choice and trip-making equations. I use the resulting parameter estimates to evaluate the counterfactuals, holding population, facility areas and calibration parameters fixed. I report 95 per cent confidence intervals based on the bootstrap distribution.

7.3 Distance Sensitivity and Trip-Making

The estimates reveal two sources of gender differences in travel behaviour. First, women are more sensitive to distance when choosing destinations. For the RPTRA activity composite, the estimated distance coefficient is 1.02 for women and 0.77 for men, implying that an additional kilometre reduces relative destination flows by 64 and 54 per cent, respectively (Table 5). The female-to-male coefficient ratio is 1.33, with a 95 per cent bootstrap interval of $[1.24, 1.38]$. These estimates are consistent with the shorter distances travelled by women for non-work activities documented in Section 5.

Table 5: Destination choice and trip-making: estimated parameters by purpose

Purpose	δ^m	δ^f	F/M	λ^m	λ^f	Trip-making M	Trip-making F
Work and business	0.346	0.373	1.08	0.011	0.119	53.15	21.52
Shopping	0.703	1.478	2.10	0.114	0.089	5.61	26.95
Broad non-work	0.730	1.149	1.57	0.020	0.089	22.86	41.50
Recreation	1.070	0.897	0.84	0.197	0.348†	1.14	0.48
RPTRA composite	0.768	1.020	1.33	0.108	0.317	3.63	5.09

Notes: Estimates use the 2002 SITRAMP survey for adults aged 15 and above across 261 mainland neighbourhoods. δ^g measures gender-specific distance sensitivity from PPML regressions of origin–destination trip flows, with gender-specific origin and common destination fixed effects. F/M is the ratio of female to male distance coefficients. λ^g is the coefficient on log access from the person-weighted trip-making logit, estimated separately by gender with household controls and sub-district fixed effects. Trip-making refers to making at least one qualifying home-originating trip, including individuals who made no trips. Broad non-work includes shopping, private and other purposes; the RPTRA composite includes recreation and private trips to educational and medical destinations. Work includes business, and purpose definitions may overlap. Appendix E reports additional specifications and estimation details.

Second, women's trip-making is more responsive to access to destinations. The estimated coefficient on log access is 0.317 for women and 0.108 for men, implying that a 10 per cent increase in access raises the odds of making a qualifying trip by approximately 3 per cent for women and 1 per cent for men. The difference between the coefficients is 0.209, with a 95 per cent bootstrap interval of $[0.024, 0.378]$ and a

test of equality yielding $p = 0.015$. Thus, improvements in access are associated with a larger increase in women's likelihood of undertaking trips for RPTRA activities.

These gender differences vary across activities. Women's distance sensitivity is substantially greater for broad non-work trips, with a female-to-male coefficient ratio of 1.57, compared with 1.08 for work and business trips and 1.07 for school trips. The larger female response of trip-making to access is less consistent across individual purposes: I find no statistically significant gender difference for shopping or broad non-work trips (Appendix E.12). The main counterfactual therefore focuses on the RPTRA composite, which captures the range of activities supported by these neighbourhood public spaces.

7.4 New Destinations and Additional Activity

I introduce RPTRA-like facilities into the set of destinations estimated from the 2002 travel diaries. The baseline retains existing opportunities, including other recreational destinations, but excludes the new RPTRA network. Adding a facility increases the set of available destinations and improves access to the modelled activity, with the magnitude of this improvement depending on distance and existing alternatives.

I assume that a facility k of area s_k has attractiveness

$$A_k = a \left(\frac{s_k}{s_0} \right)^{\rho}, \tag{8}$$

where s_0 is the median RPTRA area, a determines the attractiveness of a facility of median size, and ρ governs how attractiveness varies with area. For a network of new facilities $\mathcal{N}$, access for residents of neighbourhood i becomes

$$\Phi^g_{iq}(\mathcal{N}) = \widehat{B}^g_{iq} + \sum_{k \in \mathcal{N}} A_k e^{-\widehat{\delta}^g_q d_{ik}}, \tag{9}$$

where $\widehat{B}^g_{iq}$ measures access to existing destinations. Each new facility contributes to access according to its attractiveness, distance from the resident's neighbourhood and gender-specific distance sensitivity.

For a single facility k, the change in log access is

$$\Delta \ln \Phi^g_{iq} = \ln \left(1 + \frac{A_k e^{-\widehat{\delta}^g_q d_{ik}}}{\widehat{B}^g_{iq}} \right). \tag{10}$$

The equation shows that a facility generates a larger proportional improvement in access when it is closer to residents or when they have fewer attractive alternatives at baseline. The same facility can therefore produce different access gains across neighbourhoods and between women and men.

I specify alternative values for facility attractiveness relative to existing destinations. Under the normalisation that median positive destination attractiveness equals one, I set $a = \alpha$ and consider $\alpha \in \{0.01, 0.05, 0.10\}$. I also vary ρ to examine how the results depend on returns to facility size. Dividing a fixed area among smaller facilities increases their combined intrinsic attractiveness when $\rho < 1$, reduces it when $\rho > 1$, and leaves it unchanged when $\rho = 1$. Since the historical travel diaries do not identify these parameters for new RPTRAs, the counterfactuals are evaluated under alternative assumptions about their attractiveness and size. Table 6 summarises the model parameters, values and their sources.

Table 6: Model parameters, values, and sources

Definition	Parameter (value)	Source
Destination choice		
Distance sensitivity, women	$\delta_q^f = 1.020$	Estimated
Distance sensitivity, men	$\delta_q^m = 0.768$	Estimated
Destination attractiveness	A_{jq}	Estimated
Trip-making		
Access-response coefficient, women	$\lambda_q^f = 0.317$	Estimated
Access-response coefficient, men	$\lambda_q^m = 0.108$	Estimated
Trips per trip-maker, women	$\bar{m}_q^f = 1.13$	Calibrated
Trips per trip-maker, men	$\bar{m}_q^m = 1.08$	Calibrated
New facilities		
Facility attractiveness	$a = 0.05$ (alternatives: $0.01,\ 0.10$)	Assumed
Area elasticity	$\rho = 1$ (sensitivity: 0.5–2)	Assumed
Reference facility area	$s_0 = 1{,}559\ \mathrm{m}^2$	Calibrated, RPTRA register

Notes: Values refer to the main RPTRA activity composite. Distance sensitivity and destination attractiveness are estimated using PPML gravity regressions. Access-response coefficients are estimated using logit regressions of trip-making. Destination attractiveness is normalised to a median positive value of one. Facility attractiveness and its area elasticity are assumed, with alternative values examined in sensitivity analyses.

An improvement in access affects the probability of undertaking an activity through the trip-making equation. Let $H_{iq}^g(\mathcal{N}) = \sum_{k\in\mathcal{N}} A_k e^{-\widehat{\delta}_q^g d_{ik}}$ denote the additional access provided by the new network. Holding household characteristics fixed, the change in the log odds of making a trip is

$$\Delta \ln\left(\frac{P_{nq}^g}{1-P_{nq}^g}\right) = \lambda_q^g \ln\left(1 + \frac{H_{iq}^g(\mathcal{N})}{\widehat{B}_{iq}^g}\right). \tag{11}$$

The change in trip-making depends on the improvement in access relative to existing opportunities, the estimated access-response coefficient λ_q^g, and the individual's initial probability of making a trip. Aggregate gains also depend on the number of residents reached. A facility located in a densely populated neighbourhood may therefore generate fewer additional trips than one serving fewer residents with more limited access to existing destinations.

I use the model to distinguish new trips generated by RPTRA provision from trips redirected from existing destinations. Let $P_{iq}^g(\mathcal{N})$ denote the average probability of making a trip among residents of neighbourhood i, and let $\varnothing$ denote the absence of the added RPTRA network. The net increase in trips for activity q from a network $\mathcal{N}$ is

$$G^g(\mathcal{N}) = \bar{m}_q^g \sum_i N_i^g \left[P_{iq}^g(\mathcal{N}) - P_{iq}^g(\varnothing)\right], \qquad \sum_{k\in\mathcal{N}} s_k = \bar{S}. \tag{12}$$

Here $\bar{m}_q^g$ is the average number of trips per trip-maker, N_i^g is the adult population of neighbourhood i, fixed at its 2010 level, and $\bar{S}$ is the total land allocated to the facilities. I construct neighbourhood trip-making probabilities by averaging individual predictions using survey weights, preserving differences in household characteristics within neighbourhoods. The equation measures additional trips for activity q, rather than changes in whether individuals leave home for any purpose.

Not all trips to the new facilities are additional. Some are made by individuals who would otherwise

have visited an existing destination for the same activity. The destination-choice probabilities determine the share of trips attracted to the new facilities, while the trip-making equation determines how much total activity increases. For a given gender and purpose, let $P_i = P_{iq}^g(\varnothing)$ and $P_i' = P_{iq}^g(\mathcal{N})$ denote trip-making probabilities before and after adding the facilities, and let $B_i = \widehat{B}_{iq}^g$ and $H_i = H_{iq}^g(\mathcal{N})$ denote existing and additional access. The resulting decomposition is

$$\begin{aligned} \text{new-site trips}_i &= \bar{m} N_i P_i' \frac{H_i}{B_i + H_i}, \\ \text{additional trips}_i &= \bar{m} N_i (P_i' - P_i), \\ \text{redirected trips}_i &= \bar{m} N_i \left(P_i - P_i' \frac{B_i}{B_i + H_i} \right). \end{aligned} \tag{13}$$

Trips to the new facilities are therefore the sum of additional and redirected trips. $0 \leq \lambda_q^g \leq 1$ ensures that additional trips do not exceed new-site visits, so redirected trips remain non-negative. This decomposition captures substitution between destinations within the same activity; Appendix E.8 extends the analysis to the probability of making any trip across purposes under additional assumptions.

7.5 The Observed Network

I first use the model to predict how many additional trips the existing RPTRA network would generate relative to a baseline without these facilities. The network comprises 315 RPTRAs across 168 neighbourhoods, covering a total area of 55.78 hectares, with a median facility size of 1,559 square metres. I retain the observed distribution of facilities and their sizes, but represent their locations using the centroids of their host neighbourhoods for consistency with the model's spatial units. I set facility attractiveness to $a = 0.05$ and the area elasticity to $\rho = 1$, as in the main specification.

Under the main assumptions, the observed RPTRA network generates an estimated 4,986 additional female trips and 1,151 additional male trips per day for the composite activity, corresponding to increases of 2.51 and 0.83 per cent over baseline trip-making, respectively (Appendix Table E24). However, not all visits to the new facilities represent additional activity. The model predicts that 31 per cent of women's visits and 11 per cent of men's visits are additional, with the remainder redirected from existing destinations. The facilities therefore attract substantially more visits than the net increase in trips they generate.

The larger increase in women's trips reflects differences in both distance sensitivity and the response of trip-making to access. To assess their relative importance, I replace women's estimated parameters with the corresponding male parameters, one at a time, while holding their baseline trip-making, population and household characteristics fixed. Replacing women's distance sensitivity reduces predicted additional trips by approximately 6 per cent, from 4,986 to 4,664. Replacing their access-response coefficient reduces additional trips by 66 per cent, to 1,674 (Table 7). Thus, the larger predicted female gain arises primarily from the stronger relationship between access and trip-making, rather than differences in the distance over which destinations are attractive.

The distinction between visits and additional trips also varies by activity. Under the same facility-attractiveness assumptions, only 7.4 per cent of predicted shopping visits and 5.4 per cent of broad non-work visits by women are additional, compared with 31.0 per cent for the RPTRA composite. New destinations can therefore attract considerable activity even when most visits substitute for trips to existing

Table 7: Women's additional trips under alternative travel parameters

Women's parameters	Trips/day	Change
Baseline	4,986	–
Distance parameter replaced with men's	4,664	-6.5%
Access response replaced with men's	1,674	-66.4%

Notes: Predicted additional trips for women under the observed RPTRA network, with $a = 0.05$ and $\rho = 1$. Each specification replaces one female parameter with its male counterpart, holding baseline trip-making, population and household characteristics fixed. The comparisons quantify the contribution of each parameter within the model.

destinations (Appendix E.5).

7.6 Counterfactual Allocation of Public Spaces

I use the model to examine how the location of public spaces affects the number of additional trips generated from a fixed quantity of land. I allocate the observed RPTRA network's total area of 55.78 hectares across 100 equal-sized facilities, varying only their locations. This holds the number, size and total attractiveness of facilities constant across allocations. I compare three placement rules: locating facilities in the most densely populated neighbourhoods, prioritising neighbourhoods with low existing access, and choosing locations to maximise total additional trips, $G^f(\mathcal{N}) + G^m(\mathcal{N})$.

The population-density rule selects the 100 neighbourhoods with the most adult residents per square kilometre. The low-access rule instead prioritises neighbourhoods based on their population and existing access to destinations, using the index $N_i^m/B_i^m + N_i^f/B_i^f$. The third rule chooses locations to maximise predicted additional trips for women and men combined. All three allocations select from the same 261 neighbourhood centroids, with at most one facility per neighbourhood. For the model-based allocation, I use a greedy search followed by local improvements, obtaining the best-found allocation within this candidate set.

The choice of locations substantially affects predicted trip-making. Allocating facilities by population density generates $4{,}510$ additional trips per day across women and men, compared with $7{,}396$ when prioritising low existing access and $7{,}874$ under the model-based allocation (Table 8). The model-based allocation therefore generates approximately 75 per cent more additional trips than density targeting with the same public-space provision. These differences reflect not only how many residents are reached, but also their distances from new facilities and the destinations already available to them.

Table 8: Additional trips from alternative allocations of 100 identical facilities

Placement rule	Women	Men	Total [95% interval]
Population density	3,671	838	4,510 [2,052, 8,130]
Low existing access	5,980	1,417	7,396 [3,345, 11,969]
Total additional trips	6,378	1,496	7,874 [3,615, 12,710]

Notes: Each allocation uses 100 equal-sized facilities with a combined area of 55.78 hectares, under the main assumptions $a = 0.05$ and $\rho = 1$. Entries report net additional daily trips for the RPTRA activity composite. The population-density rule ranks adult residents per square kilometre, while the low-access rule ranks $N_i^m/B_i^m + N_i^f/B_i^f$. The model-based allocation maximises total predicted additional trips using greedy selection and local improvements. All facilities are assigned to distinct neighbourhood centroids. Brackets report 95 per cent intervals from 200 common origin-bootstrap draws, evaluating each allocation at fixed locations.

To illustrate how population and existing access jointly determine the gains from facility placement, I compare two potential locations in Jakarta: Tugu Utara, a neighbourhood with a relatively large female population and low existing access, and Gambir, a central neighbourhood with fewer female residents

and greater access to existing destinations. Tugu Utara has 24,645 adult women and a baseline female access index of 1.26, compared with 1,137 women and an index of 8.85 in Gambir. In a separate exercise allocating the entire 55.78 hectares to a single facility, the model predicts 6,596 additional female trips per day when located in Tugu Utara, compared with 1,066 in Gambir. These predictions include gains among residents across the city, not only those living in the host neighbourhood. The comparison illustrates how the benefits of a new facility depend on the population it can reach and the alternatives already available to those residents.

Women's greater estimated response to improvements in access does not necessarily imply that women and men benefit from very different facility locations. I compare allocations of the same 100 facilities that separately maximise additional female and male trips. The allocation chosen to maximise women's trips generates 6,379 additional female trips per day, compared with 6,355 under the allocation chosen for men (Table 9). The difference is only 24 trips, or 0.38 per cent, with a paired 95 per cent interval of $[-0.27, 1.20]$ per cent. Thus, under the main assumptions, locations that maximise men's additional trips also generate nearly the same gains for women, despite their different estimated responses to access.

Table 9: Gender-targeted allocations of 100 identical facilities

Allocation objective	Women	Men	Total
Men's additional trips	6,355	1,499	7,854
Women's additional trips	6,379	1,495	7,874
Female loss under male objective: 0.38%; 95% interval [-0.27, 1.20]			

Notes: Each allocation uses 100 equal-sized facilities and 55.78 hectares of land, under the main assumptions $a = 0.05$ and $\rho = 1$. Locations are chosen separately to maximise women's or men's additional trips and are evaluated for both genders. Brackets report the paired 95 per cent interval for the percentage difference in women's additional trips between the two allocations. The results concern additional trips for the RPTRA activity composite.

7.7 Discussion

The counterfactuals show that where public spaces are located matters for how many additional trips they generate. This finding also holds when I consider allocations involving one, three or six larger parks and networks of 315 smaller facilities (Appendix E.15). Comparisons across different facility sizes depend on the assumed relationship between area and attractiveness, while the main exercise with 100 identical facilities holds size and total land fixed to isolate the effect of location. The quantitative gains also depend on the assumed attractiveness of new facilities, which I vary in Appendix E.5.

The model provides a different perspective on the reduced-form results. The reduced form estimates measure changes in the number of people visible near RPTRAs, whereas the model predicts how access to new destinations changes trip-making. It also distinguishes additional trips from visits redirected from existing destinations. These measures need not move together because street presence reflects not only journeys but also how long people remain visible, whether visiting a facility or undertaking other activities nearby. This helps explain why the larger predicted gains in women's trips can coexist with similar proportional changes in female and male street presence.

The model also offers a general lesson for the allocation of urban amenities. A new facility generates more activity when it serves residents with limited existing access, but the number of people reached and their responses to improved access also matter. Consequently, locating facilities in the most densely

populated neighbourhoods need not maximise additional trips. The analysis holds residences, employment and existing destinations fixed to isolate this mechanism, and does not account for land costs, congestion or longer-run adjustments. The model shows that allocating public spaces based on population density alone can miss substantial gains in trip-making, particularly where residents have limited access to existing destinations.

8 Conclusion

Cities determine not only which opportunities are available, but also where they are located and who is able to reach and use them. This paper examines gender differences in the use of urban public space in Jakarta. Using repeated street imagery, I construct a new measure of street presence that captures a dimension of urban activity distinct from employment, residential location and trip-making. Women account for approximately one-quarter of identified adults on Jakarta's streets, with roughly three men observed for every woman. This gender gap persists over time and varies considerably across neighbourhoods.

Differences in employment and travel help account for the gender gap, but do not fully explain the much larger disparity in street presence. Under the baseline assumptions about workplace visibility, employment, working hours, work location and commuting account for at most 69.8 per cent of the observed gap. Travel patterns reveal a further distinction: women make approximately 82 per cent as many trips per adult as men, despite a female-to-male street-presence ratio of only 33 per cent. Women also make more non-work trips and travel shorter distances for these activities. Together, these findings suggest that gender differences in street presence reflect not only how frequently women and men leave home, but also where they go, the activities they undertake and how long they remain visible in public.

The introduction of neighbourhood parks provides a setting to examine how urban amenities change these patterns. I find that RPTRA openings increase pedestrian presence on streets immediately around their boundaries for both women and men, without a statistically distinguishable difference in their proportional responses. Yet women constitute a substantially higher share of adults photographed inside RPTRAs than on surrounding streets. Women are also particularly under-represented among people observed sitting on streets, while men's street presence more frequently involves activities such as vending. These patterns suggest that the use of a public-space amenity and the number of people visible around it need not move together. A facility may attract women who spend much of their time inside, while also attracting street-based activities that keep other users visible for longer.

To examine how access to new destinations shapes urban activity, I develop a model of destination choice and trip-making. The model distinguishes additional trips generated by new facilities from visits redirected from existing destinations. Under the main model assumptions, 31 per cent of predicted female visits and 11 per cent of male visits to the added RPTRA network represent additional trips for the modelled activities. Women's larger predicted gains arise primarily from their stronger estimated relationship between access and trip-making. The model also shows that where public spaces are located matters substantially. Holding the number, size and total land area of 100 facilities fixed, the model-based allocation generates approximately 75 per cent more additional trips than population-density targeting. This difference arises because the benefits of a new destination depend on the population it reaches and the alternatives those residents already have.

These findings have implications for how cities evaluate and provide public spaces. Increasing local street activity, generating additional trips and narrowing gender gaps in public-space use are distinct

policy objectives. A facility may attract substantial female use without proportionately changing women's street presence, while locations that attract many visitors need not generate the most additional trips. The model provides a way to incorporate differences in population, travel costs and existing access into facility-placement decisions, while holding residential locations and other urban opportunities fixed. The model evaluates additional trip-making under a common land constraint, rather than monetary welfare.

The findings also raise questions about both the determinants and consequences of public-space use. The measurement approach developed in this paper provides a way to study how different groups occupy urban space across neighbourhoods and over time, and how their presence changes following investments in transport, safety and urban amenities. One direction for future research is to distinguish constraints on women's use of the city from differences in preferences. Lower female street presence may reflect constraints related to safety, gender norms and household responsibilities, alongside differences in preferred activities and destinations. Understanding these mechanisms would help establish whether expanding access alone is sufficient to narrow gender gaps in public-space use.

A second direction is to examine whether changes in public-space use affect other economic and social outcomes. The presence of more people, and changes in their gender composition, may alter the experience of public space for other users. Greater activity can increase demand for local businesses, while changes in who occupies streets may influence perceived safety, social interactions and norms around women's presence in public. These effects could, in turn, shape women's access to employment, residential choices and neighbourhood development, including possible gentrification or displacement. Increased activity may also create costs through congestion or competition for space. Establishing these effects requires linking changes in street presence to subsequent outcomes and distinguishing the role of public-space activity from the interventions that generate it. Combining repeated imagery with data on journeys, time spent at destinations and other economic outcomes offers a way to investigate these questions. More broadly, studying who is present in the city, rather than only where people live or work, can help establish how the benefits of urban development are distributed across different groups.

References

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf, "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 2015, *83* (6), 2127–2189.

Badan Pusat Statistik, "Population Census: Jakarta Neighbourhood Population Tabulations," 2010. Adult population counts by gender used in the counterfactuals.

_ , "Potensi Desa (PODES)," Badan Pusat Statistik, Jakarta 2011–2024. Village census; waves used are 2011, 2014, 2018 and 2019.

_ , "Statistik Komuter Jabodetabek," Technical Report, Badan Pusat Statistik, Jakarta 2014.

_ , "Sensus Ekonomi 2016 (SE2016)," Badan Pusat Statistik, Jakarta 2017. Establishment and employment counts; kelurahan aggregates.

_ , "Statistik Komuter Jabodetabek," Technical Report, Badan Pusat Statistik, Jakarta 2019.

_ , "Statistik Komuter Jabodetabek," Technical Report, Badan Pusat Statistik, Jakarta 2023.

Badan Pusat Statistik Provinsi DKI Jakarta, "Jumlah Penduduk Hasil SP2020 Provinsi DKI Jakarta sebesar 10.56 juta jiwa," Press release, 22 January 2021.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Gender Differences in Job Search: Trading off Commute against Wage," *Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Berita Jakarta, "DKI Bangun 100 RPTRA di 2017," Official Jakarta provincial news service, 17 December 2016.

Borker, Girija, "Safety First: Perceived Risk of Street Harassment and Educational Choices of Women," Policy Research Working Paper 9731, The World Bank 2021.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Brier, Glenn W., "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review*, 1950, *78* (1), 1–3.

Butts, Kyle, "Difference-in-Differences Estimation with Spatial Spillovers," Technical Report 2105.03737, arXiv 2023.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Clarke, Damian, "Estimating Difference-in-Differences in the Presence of Spillovers," MPRA Paper 81604, University Library of Munich 2017.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

de Madariaga, Inés Sánchez, "Mobility of Care: Introducing New Concepts in Urban Transport," in Inés Sánchez de Madariaga and Marion Roberts, eds., *Fair Shared Cities: The Impact of Gender Planning in Europe*, Routledge, 2013, pp. 33–48.

_ **and Elena Zucchini**, "Measuring Mobilities of Care, a Challenge for Transport Agendas: From One to Many Tracks," in Inés Sánchez de Madariaga and Michael Neuman, eds., *Integrating Gender into Transport Planning: From One to Many Tracks*, Palgrave Macmillan, 2019, pp. 145–173.

Deng, Yubin, Ping Luo, Chen Change Loy, and Xiaoou Tang, "Pedestrian Attribute Recognition at Far Distance," in "Proceedings of the ACM International Conference on Multimedia" 2014, pp. 789–792.

Donaldson, Dave and Adam Storeygard, "The View from Above: Applications of Satellite Data in Economics," *Journal of Economic Perspectives*, 2016, *30* (4), 171–198.

Garlick, Robert, Erica Field, and Kate Vyborny, "Women's Mobility and Labor Supply: Experimental Evidence from Pakistan," IZA Discussion Paper 17883, Institute of Labor Economics (IZA) 2025. Also World Bank Policy Research Working Paper 11256. Author order is randomised.

Gauvin, Laetitia, Michele Tizzoni, Simone Piaggesi, Andrew Young, Natalia Adler, Stefaan Verhulst, Leo Ferres, and Ciro Cattuto, "Gender Gaps in Urban Mobility," *Humanities and Social Sciences Communications*, 2020, *7*, 11.

Gebru, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei, "Using Deep Learning and Google Street View to Estimate the Demographic Makeup of Neighborhoods across the United States," *Proceedings of the National Academy of Sciences*, 2017, *114* (50), 13108–13113.

Glaeser, Edward L., Scott Duke Kominers, Michael Luca, and Nikhil Naik, "Big Data and Big Cities: The Promises and Limitations of Improved Measures of Urban Life," *Economic Inquiry*, 2018, *56* (1), 114–137.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Google, "Google Places API," https://developers.google.com/maps/documentation/places 2025. Amenity locations and categories, queried 2025.

_ , "Google Street View Static API," https://developers.google.com/maps/documentation/streetview 2025. Panorama imagery and metadata for DKI Jakarta, 2013–2024.

Harari, Mariaflavia and Maisy Wong, "Slum Upgrading and Long-Run Urban Development: Evidence from Indonesia," *Review of Economic Studies*, 2026, *93* (4), 2646–2679.

ICF International, "Demographic and Health Surveys: Indonesia 2017," ICF, Rockville, MD 2018.

Japan International Cooperation Agency, "The Study on Integrated Transportation Master Plan for Jabodetabek (SITRAMP)," Technical Report, Japan International Cooperation Agency and National Development Planning Agency, Jakarta 2002.

Karekurve-Ramachandra, Varun and Gaurav Sood, "Missing Women on Indian Streets," Technical Report, SocArXiv 2026.

Kementerian Agraria dan Tata Ruang / Badan Pertanahan Nasional, "Peta Zona Nilai Tanah (ZNT), Provinsi DKI Jakarta," Badan Pertanahan Nasional, Jakarta 2015, 2019. Assessed land value per square metre by valuation zone.

Kleven, Henrik and Camille Landais, "Gender Inequality and Economic Development: Fertility, Education and Norms," *Economica*, 2017, *84* (334), 180–209.

Kreindler, Gabriel, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A. Olken, "Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta," *American Economic Review*, 2026, *116* (9), 3330–3379.

Kreindler, Gabriel E. and Yuhei Miyauchi, "Measuring Commuting and Economic Activity inside Cities with Cell Phone Records," *Review of Economics and Statistics*, 2023, *105* (4), 899–909.

Kull, Meelis, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach, "Beyond Temperature Scaling: Obtaining Well-Calibrated Multi-Class Probabilities with Dirichlet Calibration," in "Advances in Neural Information Processing Systems," Vol. 32 2019.

Liu, Xihui, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Shuai Yi, Junjie Yan, and Xiaogang Wang, "HydraPlus-Net: Attentive Deep Features for Pedestrian Analysis," in "Proceedings of the IEEE International Conference on Computer Vision" 2017, pp. 350–359.

Lloyd's Register Foundation, "World Risk Poll 2023," Lloyd's Register Foundation, London 2023.

McFadden, Daniel, "The Measurement of Urban Travel Demand," *Journal of Public Economics*, 1974, *3* (4), 303–328.

Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J. Redding, "The Economics of Spatial Mobility: Theory and Evidence Using Smartphone Data," *Quarterly Journal of Economics*, 2025, *140* (4), 2507–2570.

Ng, Wei-Shiuen and Ashley Acker, "Understanding Urban Travel Behaviour by Gender for Efficient and Equitable Transport Policies," Discussion Paper 2018-01, International Transport Forum, OECD 2018.

Olivetti, Claudia, Jessica Pan, and Barbara Petrongolo, "The Evolution of Gender in the Labor Market," in "Handbook of Labor Economics," Vol. 5, Elsevier, 2024, chapter 8, pp. 619–677.

OpenStreetMap contributors, "OpenStreetMap," https://www.openstreetmap.org 2025. Historical amenity extracts, 2013–2025.

Pemerintah Provinsi DKI Jakarta, "Peraturan Gubernur Nomor 196 Tahun 2015 tentang Pedoman Pengelolaan Ruang Publik Terpadu Ramah Anak," 2015.

_ , "Peraturan Gubernur Nomor 139 Tahun 2016 tentang Pemanfaatan Ruang untuk Kegiatan Ruang Publik Terpadu Ramah Anak," 2016. Article 4 and Annex I.

_ , "Daftar Ruang Publik Terpadu Ramah Anak (RPTRA)," Dinas Pemberdayaan, Perlindungan Anak dan Pengendalian Penduduk, Jakarta 2021. Site names, addresses, coordinates, boundary polygons and opening dates.

Redding, Stephen J. and Esteban Rossi-Hansberg, "Quantitative Spatial Economics," *Annual Review of Economics*, 2017, *9*, 21–58.

Sant'Anna, Pedro H. C. and Jun Zhao, "Doubly Robust Difference-in-Differences Estimators," *Journal of Econometrics*, 2020, *219* (1), 101–122.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, 2006, *88* (4), 641–658.

Wooldridge, Jeffrey M., "Simple Approaches to Nonlinear Difference-in-Differences with Panel Data," *The Econometrics Journal*, 2023, *26* (3), C31–C66.

A Measurement

This appendix documents the street-imagery samples, their weighting, classification and validation. It then reports the supplementary descriptive evidence cited in Sections 3 to 5: the geography and persistence of the gap, observed street activity, time use and safety, and cross-country context. None of these exercises identifies the effects of RPTRA openings.

Table A1: Data sources, coverage and measured outcomes

Source	Period	Role
Street imagery	Machine: 2013–2019; hand labels through 2024	Pedestrian presence, spatial variation and opening comparisons.
SITRAMP travel diaries	2002	Full-day descriptions and home-departure model, with distinct restrictions.
SAKERNAS	2018, 2023	Employment, hours and work-location visibility assignments.
BPS commuter surveys	2014, 2019, 2023	Commuting patterns; walking-exposure benchmark.
Population census	2010	Fixed model populations and neighbourhood composition.
Economic Census	2016	Employment geography.
RPTRA register and boundaries	Openings from 2015	Dates, locations and area of 315 facilities.
Interior photographs	Selected user uploads	Hand-labelled composition inside facilities.
Population register, PODES, assessed land values	Source-specific waves	Context, residential response and supplementary outcomes.

Notes: Street imagery records pedestrians visible on covered streets; travel diaries record journeys by surveyed residents; labour-force surveys describe resident work activity; and the RPTRA register records facilities and opening dates. Model populations are adults aged 15 and over in the 2010 neighbourhood data. Neighbourhood comparisons use the 261 mainland kelurahan. The text and appendices describe source-specific samples, weights and restrictions.

A.1 Sampling and Collection Protocol for the Hand-Labelled Sample

The sampling frame contains 63,136 100-metre grid cells with centroids inside Jakarta's administrative boundary, excluding the Thousand Islands. A ten-per-cent random sample gives 6,314 candidate locations. Collection stops at a fixed labelling quota rather than covering every candidate. The final sample contains 10,216 panoramas at 2,963 locations, averaging 3.45 panoramas per location. The retained locations represent 4.7 per cent of the frame. Locations without eligible imagery do not enter the sample.

Search rule. When a sampled centroid has no Street View coverage, the enumerator selects the nearest eligible Google-sourced panorama within the same block. If none is available, the location is excluded rather than matched across a road boundary. Private contributions are excluded because their capture conditions and viewpoints differ from vehicle-mounted imagery. For coverage inside a bus or metro station, the enumerator uses the nearest street outside. For coverage on a flyover, the street beneath is used.

Viewing rule. The viewing protocol fixes the orientation to north, disables navigation arrows and prohibits zoom. Each location-period observation therefore comes from one panorama, viewed under the same rules across enumerators, rather than from several adjacent captures.

Period selection. Panoramas from 2020, 2021 and 2025 are excluded. When a location has no more than two eligible capture periods, both are retained where available. Otherwise, the first and last are retained with intermediate captures separated by at least 24 months. This limits closely spaced repeat observations without imposing a common capture month.

Labelling scheme. The protocol records people in public areas outside buildings and gated premises, including roads, sidewalks, gazebos, parks and shaded stalls. It excludes private terraces and balconies and people inside moving motorised vehicles. Non-motorised riders are recorded separately and excluded from the pedestrian-only accounting sample in Section C.4. Each person is classified as an infant or toddler, child, or adult; adults are further classified as male, female or of unidentifiable gender. Adult labels also record posture, companionship, women's head coverings and whether they are buying, selling or street vending. The full instrument and its consistency-check flowchart are available on request.

Physical environment. Enumerators also record road type and lane count; the presence and continuity of pavement; vehicle types; and trees, bushes, benches, lighting, encroachment, transit stops, refuse and congestion. Capture month and year come from the imagery. Shadow direction and length provide approximate categories for early morning, late morning, midday, afternoon and evening. Images whose shadows cannot be interpreted are assigned to the overcast category.

A.2 Machine-Panel Coverage and Weights

The machine-panel weights combine the probability that a cell is sampled with the probability that it returns usable imagery. Let s_i be the cell-inclusion probability under the four-phase design and let $\widehat{q}_i$ be its fitted coverage probability conditional on being sampled. The weight is

$$w_i = \frac{1}{s_i \widehat{q}_i \widehat{d}_i \widehat{r}_i}.$$

The sampling probability is known and varies by distance stratum. The coverage probability is estimated on the full sampled frame, including cells without usable imagery. Its inverse is the coverage component. The production weight additionally divides by centroid-level download retention $\widehat{d}_i$ and complete-panorama segmentation retention $\widehat{r}_i$. These retention rates are computed within sampling-phase by distance-stratum cells, with a stratum-level fallback for cells below 200 centroids; they are not annual image-level response rates.

Road accessibility affects which locations can enter the imagery. The weights reweight observed cells for estimated coverage differences, but do not recover activity in places that Street View never observes. Nor do cell fixed effects or coverage weights alone rule out time-varying selection around an opening. The estimand remains tied to the covered street network.

Figure A1 shows the machine-panel sample locations and an example of Street View coverage in a residential area of North Jakarta.

Figure A1: The street-image sample and Street View coverage



(a) Sample locations **(b)** Street View coverage **(c)** Sample cells and panoramas

Notes: Panel (a) maps machine-panel panorama locations in mainland Jakarta, 2013–2019. Panels (b) and (c) show the same residential area of Sunter, North Jakarta. In panel (b), highlighted lines mark Google Street View coverage. Panel (c) overlays the 50-metre-wide hexagonal sampling grid on OpenStreetMap. Large points are sampled cell centroids, coloured by whether imagery was found; small points mark available panoramas. Main roads have imagery, while many narrow lanes do not. Sampling and coverage weights reweight observed locations but cannot recover activity on streets without imagery. Map data: Google (panel b); OpenStreetMap contributors (panel c).

A.3 Street Occupancy and Aggregate Mobility Flows

I compare street occupancy with the released Veraset origin–destination data in the replication package for *Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta*. The package's README describes confidential smartphone records aggregated to origin–destination–time cells. This comparison uses the released aggregates and geography, not individual trajectories.

The published flows are population-expanded, daily-equivalent weekday trips. I retain the 51 weeks wholly within 2019 using the supplied date crosswalk, average across those weeks and sum arrivals and departures on the package's one-kilometre hexagonal grid. A trip remains included when its other endpoint lies outside the Street View-covered area. Within-grid trips contribute once to arrivals and once to departures. Dividing by grid area gives trip density. The morning-arrivals measure uses trips starting between 05:00 and 10:00, following the source code.

The street outcome is the sampling-weighted mean of expected women plus men per panorama in 2019. I assign street-cell centroids to the same grid. Of 47,641 street cells, one has no grid match and is recorded separately. The comparison uses 770 grids with both imagery and mobility records; missing flows are not coded as zero. Street intensity remains a mean per sampled panorama, not an expanded number of distinct people.

Across the matched grids, the Pearson correlation of log-one-plus arrival density with log-one-plus street intensity is 0.458; the rank correlation is 0.467. Using both trip ends gives correlations of 0.454 and 0.457. Restricting the sample to grids with at least 20 street cells gives similar results (Table A2 and Figure A2). The two sources therefore show positive spatial correspondence in aggregate activity, without establishing that either captures all activity.

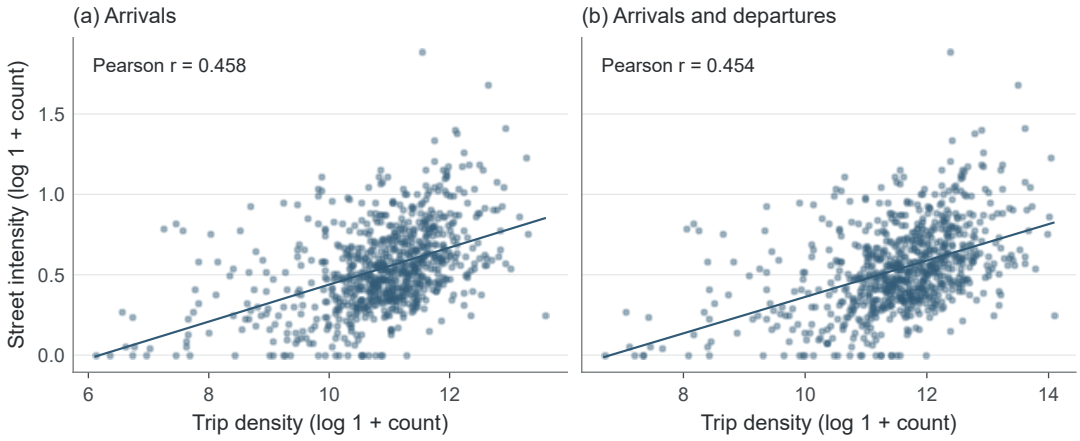
This comparison does not validate gender shares or model counterfactuals. The mobility extract has no gender, purpose, route exposure or destination dwell time. Smartphone selection also differs from camera coverage. A busy enclosed destination can produce many recorded trips while placing few people in street view, so the two measures need not be proportional.

Table A2: Spatial correlations between street presence and aggregate trip flows, 2019

Flow measure	Grids	Pearson	Spearman
All matched grids			
Arrivals	770	0.458	0.467
Departures	770	0.448	0.447
Arrivals + departures	770	0.454	0.457
Morning arrivals	770	0.395	0.375
At least 20 street cells			
Arrivals	662	0.455	0.453
Departures	662	0.437	0.428
Arrivals + departures	662	0.447	0.440
Morning arrivals	662	0.360	0.346

Notes: 2019 Jakarta Google Street View and released Veraset aggregate flows. Each observation is a matched 1-kilometre hexagonal grid; 770 grids match, with the minimum-cell restriction shown in the table. Street intensity is the sampling-weighted mean expected adult count (women plus men) per panorama. Flows are population-expanded daily-equivalent weekday trips, averaged across 51 weeks wholly in 2019 and divided by grid area in square kilometres. Arrivals count destination endpoints; trip ends sum arrivals and departures. Morning arrivals use trip start times 05:00–10:00. Pearson correlations use log-one-plus values, Spearman correlations ranks; grids receive equal weight. No independence-based significance tests are reported.

Figure A2: Street occupancy and aggregate arrival and trip-end densities, 2019



Notes: Each point is one of 770 matched 1-kilometre hexagonal grids in Jakarta in 2019. Street intensity is the sampling-weighted mean expected adult count (women plus men) per Google Street View panorama. Veraset flows are released population-expanded daily-equivalent weekday trips, averaged over the 51 weeks wholly in 2019 and divided by grid area in square kilometres. Arrivals count destination endpoints; trip ends add arrivals and departures, including both endpoints of within-grid trips. Both axes use log-one-plus transformations. The unadjusted spatial comparison has no confidence bands and does not validate gender shares or destination dwell time.

A.4 Classifier Training, Calibration and Validation

The machine panel requires a detector to find people and an attribute classifier to describe the retained detections. A separate classifier measures visible streetscape features. The pedestrian system is trained on public pedestrian-attribute datasets, then fine-tuned on independently labelled Jakarta imagery. A separate crop audit supplies the calibration sample.

A.4.1 Training Data

PA-100K contains 100,000 pedestrian images from 598 outdoor surveillance cameras, with gender, age, clothing and accessory attributes (Liu et al., 2017). PETA combines 19,000 images from ten surveillance systems and distinguishes age groups in more detail (Deng et al., 2014). These datasets provide

pre-training. Fine-tuning uses a random sample of Jakarta Street View crops labelled independently as woman, man, child or unidentified.

A.4.2 Calibration Procedure

I use Dirichlet calibration to map the classifier's raw scores into probabilities over woman, man and child (Kull et al., 2019). The probabilities sum to one for each eligible detection. Summing them across detections gives the main probability-based counts. Hard-label rules and isotonic calibration provide alternative class assignments in the robustness checks.

Weighted crop-level cross-validation reduces the multiclass Brier score from 0.460 to 0.278 (Brier, 1950). This measures the correspondence between class probabilities and audit labels among detected crops. It does not measure the fraction of eligible pedestrians detected, or establish that detection error is stable across locations, years or crowding conditions.

A.4.3 Validation Against Hand Labels

The crop audit evaluates how detected people are classified. A comparison of complete panorama counts with human labels would additionally evaluate which people the detector misses or incorrectly retains. Calibration addresses the first question, not the second. The hand-labelled sample remains the reference for headline levels.

The hand-labelled and machine samples share broad geographic and temporal patterns. Both place East Jakarta above Central Jakarta in female share (Figure 3, panel a), but agreement at that level does not establish accurate rankings for all neighbourhoods. Table 1 summarises the validation samples and statistics, and panel (b) of Figure 3 reports the correlations of count errors with the opening design. Figure D12 compares RPTRA estimates under calibrated probabilities, isotonic probabilities and hard labels. Each gives positive boundary estimates for women and men, although hard female labels produce substantially more Poisson separation and hence different estimation support. Figure D13 compares the calibrated probability distributions.

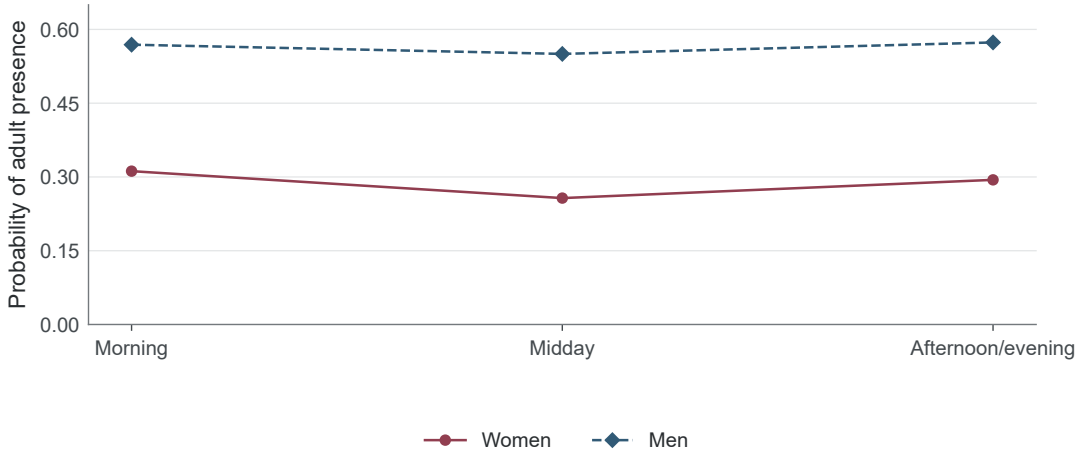
These checks leave treatment-related detection unresolved. An opening may change crowding, occlusion or the view of the site, affecting which people are detected as well as how many are present. Agreement across class-assignment rules cannot establish that detection is stable. Section D.6 discusses this limitation for the opening estimates.

A.5 Measurement Checks and Cross-City Comparisons

Time of day. Figure A3 compares hand-labelled binary presence across three assignable daylight categories based on shadow direction and length. Male presence ranges from 55.0 to 57.4 per cent and female presence from 25.7 to 31.2 per cent across morning, midday and afternoon/evening. Women are less likely to appear in each category. These are probabilities of observing at least one adult, not gender shares of person counts.[16]

[16] Unassignable captures are excluded and night-time activity is not observed. The reported concerns about travel after dark in Section A.10 do not establish the full-day occupancy gap.

Figure A3: Probability of adult presence by gender and time of day



Notes: Hand-labelled Jakarta Street View panoramas, 2013–2024 excluding 2020–2021, restricted to images with an assignable daylight category based on shadow direction and length. Each point is the unweighted proportion of panoramas with at least one identified adult of the indicated gender, using the hand-label adult categories, including separately coded non-motorised riders. A panorama can contribute to both gender indicators. These are binary presence probabilities, not shares of person counts; unassignable captures are excluded. No confidence intervals are shown, and the comparison does not describe night-time presence.

Other cities. The common coding protocol places Jakarta between Delhi and Nairobi in female representation and below London and New York (Table A3). External samples are small, coverage and capture years differ, and the unidentified-adult share varies across cities. These comparisons are therefore illustrative rather than inferential rankings. Karekurve-Ramachandra and Sood (2026) provide separate evidence of substantial gender gaps using GPS-linked wearable cameras and randomised street sampling in Mumbai and Delhi.

Table A3: Female and male share of adults observed on the street, by city

City	% adult women	% adult men	% unidentified	% of identified
Delhi	19.0	81.0	–	19.0
Bangkok	21.0	67.0	12.0	23.9
Jakarta	22.5	68.8	8.7	24.7
Nairobi	28.0	69.0	3.0	28.9
London	29.0	49.0	22.0	37.2
New York	35.0	46.0	19.0	43.2

Notes: Hand-labelled Google Street View observations from randomly sampled locations under a common coding protocol. The first three columns divide women, men and unidentified adults by all coded adults; the final column divides women by women plus men. All entries are percentages of summed counts, including separately coded non-motorised riders and excluding children; there is no population adjustment. A dash denotes an unavailable unidentified share, not a verified zero. Jakarta uses 2013–2024 excluding 2020–2021; capture years and sample sizes outside Jakarta are not harmonised in this comparison. Small external samples and differences in identifiability make it illustrative; no inferential ranking is reported.

A.6 Additional Evidence on the Gender Gap

The following exhibits describe the spatial and temporal variation in street occupancy and the activities recorded on the street. The survey evidence then provides context on time use and reported safety. These comparisons retain their source-specific populations, periods and denominators.

A.7 Additional Spatial Descriptives

Table A4: Distribution of the female share of adults observed, across neighbourhoods

Panel A: distribution across neighbourhoods	
Neighbourhoods	261
Mean	24.0
Standard deviation	2.4
10th percentile	20.9
Median	24.1
90th percentile	26.9
Minimum	18.1
Maximum	31.8
Panel B: spatial clustering	
Moran's I, 5 nearest neighbours	0.295
randomisation p	0.001
Moran's I, 8 nearest neighbours	0.249
randomisation p	0.001

Notes: Google Street View, 2013–2019, covering 261 mainland Jakarta neighbourhoods. Female adult shares, in per cent, are sampling- and coverage-weighted cell-year averages of calibrated female counts divided by female plus male counts. Neighbourhoods receive equal weight in the distribution statistics; levels depend on classifier calibration. Spatial correlation uses row-standardised five- and eight-nearest-neighbour weights, with permutation inference from 999 permutations. These descriptive statistics do not count unique residents.

A.8 Persistence of the Gap Over Time

Figure A5 reports the probability of observing at least one classified adult woman or man in the machine panel, by capture year from 2013 to 2019. Male presence is higher in every year and somewhat lower toward the end of the period; female presence is broadly stable. The plotted narrowing therefore comes mainly from the male series. Because the locations observed can change across years, the series does not by itself separate changing occupancy from changing sample composition. The binary trend uses the production $t062$ hard-label indicator: at least one retained detection assigned to the relevant gender under the 0.62 rule. It is not defined by a positive sum of calibrated probabilities. Suspect cell-years are excluded and the series uses the production observation weight.

The hand-labelled sample shows a similar but noisier pattern over the same window. Over 2013–2024, the female share first rises and then returns close to its earlier level. A change in travel modes could affect visibility because the protocol excludes people inside moving motorised vehicles, but these series do not identify such a mechanism.

A.9 Observed Activity and Its Denominators

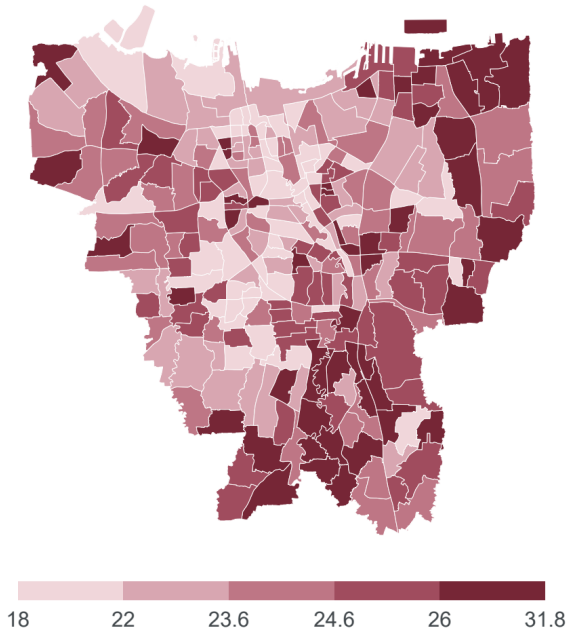
Figure A6, discussed in Section 6.4, compares posture, companionship and activity in the hand labels. Posture and activity shares use identified adults of each gender, including non-motorised riders. Companionship uses adults on foot or seated, and the final panel uses only adults recorded as trading. Accompaniment by children and street vending can overlap other categories. The figure describes people

Table A5: Employment, work location, journey composition and non-market time, by gender

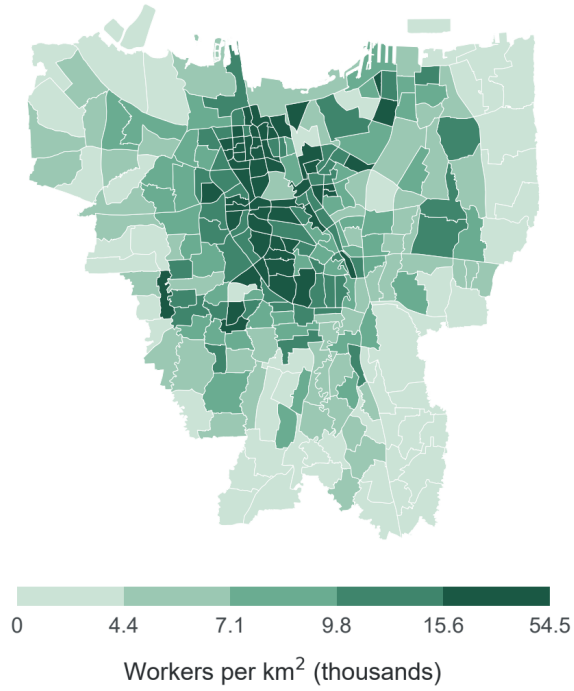
	Men	Women	Ratio	Source
A. Labour supply				
Employment-to-population ratio, adults 15+ (2018)	0.727	0.453	1.604	SAKERNAS Aug 2018
Usual weekly hours, main job (2018)	48.564	45.023	1.079	SAKERNAS Aug 2018
Employment-to-population ratio, adults 15+ (2023)	0.740	0.478	1.547	SAKERNAS Aug 2023
Usual weekly hours, main job (2023)	47.231	41.470	1.139	SAKERNAS Aug 2023
Employment on a main-activity definition (census-comparable)	0.711	0.408	1.744	SAKERNAS Aug 2018
Employed only via the probe questions, share of employed	0.032	0.055	0.582	SAKERNAS Aug 2023
B. Where work happens				
Works at own home	0.133	0.275	0.483	SAKERNAS Aug 2023
Works at an employer or client home	0.074	0.117	0.632	SAKERNAS Aug 2023
Works in an office, factory or building	0.449	0.363	1.237	SAKERNAS Aug 2023
Works in a mall or covered market	0.123	0.139	0.887	SAKERNAS Aug 2023
Works on the street, roadside or itinerant	0.163	0.064	2.561	SAKERNAS Aug 2023
Of street workers, share in machine-operator occupations	0.546	0.068	8.057	SAKERNAS Aug 2023
Home-based work, share of employed (2018)	0.125	0.252	0.496	SAKERNAS Aug 2018
Home-based work, share of employed (2023)	0.133	0.275	0.483	SAKERNAS Aug 2023
C. Commuting				
Commutes daily, share of employed (2018)	0.855	0.703	1.216	SAKERNAS Aug 2018
Walks to work, share of daily commuters (2018)	0.126	0.219	0.574	SAKERNAS Aug 2018
Commutes daily, share of employed (2023)	0.834	0.684	1.221	SAKERNAS Aug 2023
Walks to work, share of daily commuters (2023)	0.127	0.235	0.542	SAKERNAS Aug 2023
Walk-only journey, share of work commuters	0.128	0.257	0.499	Commuter survey 2023
Street-exposed transit journey, share of commuters	0.037	0.130	0.287	Commuter survey 2023
Ojek, ride-hail, taxi or shuttle, share of commuters	0.012	0.083	0.147	Commuter survey 2023
Multi-leg journey, share of commuters	0.018	0.072	0.255	Commuter survey 2023
One-way walking journey, minutes	17.658	15.583	1.133	Commuter survey 2023
Nearest transport stop within 500 m of the dwelling	0.484	0.578	0.837	Commuter survey 2023
Nearest transport stop within 1 km of the dwelling	0.787	0.859	0.916	Commuter survey 2023
D. Non-work activity				
Did household work last week	0.464	0.868	0.535	SAKERNAS Aug 2018
Did other activity last week	0.683	0.570	1.198	SAKERNAS Aug 2018
Household work took the most time last week	0.065	0.446	0.146	SAKERNAS Aug 2018
Community-activity response, mean survey code	1.569	1.702	0.922	Commuter survey 2023
E. Street presence				
Visible adults on foot or seated, 2015–19	8397.000	3025.000	2.776	Hand-labelled imagery
Adults whose gender could not be identified, share of adults, 2015–19	0.083			Hand-labelled imagery
Visible adults on foot or seated, 2022–24	6125.000	1815.000	3.375	Hand-labelled imagery
Adults whose gender could not be identified, share of adults, 2022–24	0.097			Hand-labelled imagery
Visible adults, early morning	1244.000	504.000	2.468	Hand-labelled imagery
Visible adults, late afternoon	1562.000	471.000	3.316	Hand-labelled imagery

Notes: Ratios divide men by women. Survey estimates use supplied weights for mainland DKI Jakarta; SAKERNAS includes adults aged 15 and over in the August rounds named in each row. Employment shares use all adults; hours and work-location shares condition on employment, except the explicitly labelled street-worker share. Main-activity employment is the activity occupying the most time, whereas headline employment includes probe questions. Commuter-survey rows use DKI residents with a completed travel block; each row states its commuting denominator. Street-exposed transit includes angkot, bus, train, rapid transit and multi-leg travel, excluding ojek, ride-hail, taxi, shuttle and becak. Panel D reports activity indicators except the community-activity row, which is a mean of the saved survey response code, not a participation rate. Panel E gives unweighted hand counts on foot or seated and excludes non-motorised riders; unidentified-adult shares use all coded adults. No sampling intervals are reported.

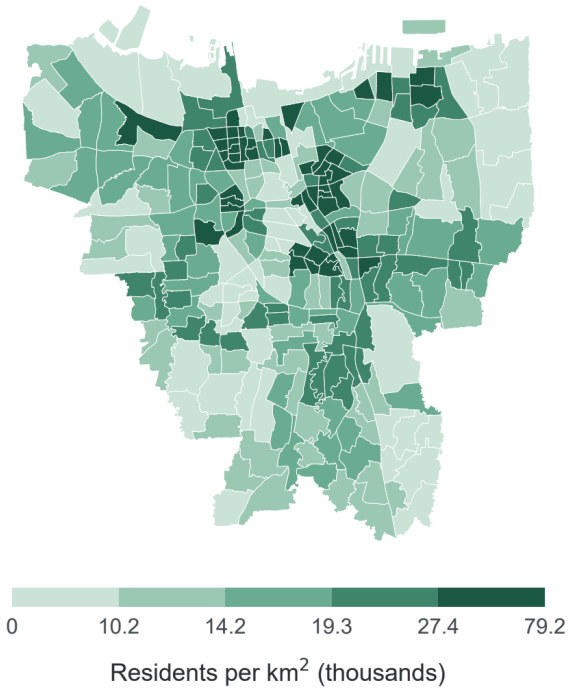
Figure A4: The female share of adults observed on the street, and the distribution of employment and population, by neighbourhood



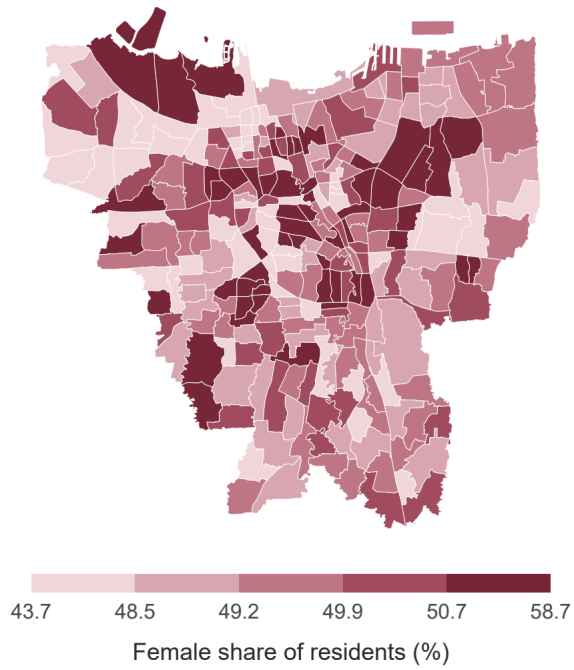
(a) Female share of adults observed



(b) Employment density



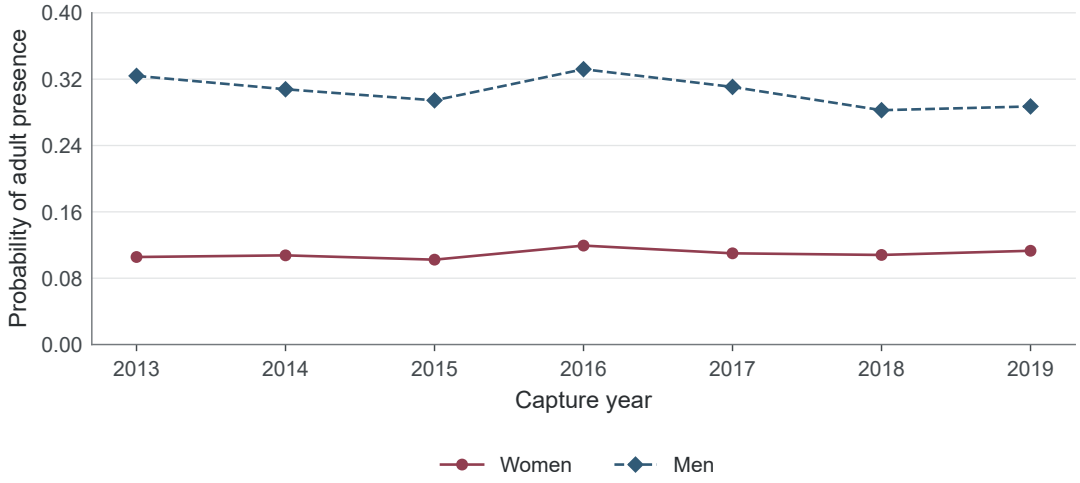
(c) Population density



(d) Female share of resident population

Notes: 261 mainland Jakarta neighbourhoods; each panel uses its own quintile bins. Panel (a) is the 2013–2019 machine-classified female share of street adults, averaging cell-year female/(female+male) calibrated counts with sampling and coverage weights. Panel (b) is 2016 Economic Census employment per square kilometre, in thousands. Panel (c) is survey-weighted 2010 Census population per square kilometre, in thousands; panel (d) is women as a percentage of all residents in that census sample. Areas use administrative neighbourhood boundaries. These are descriptive spatial distributions; street users need not reside where they are photographed, and no uncertainty shading is shown.

Figure A5: Probability of observing an adult woman and an adult man, 2013–2019



Notes: Google Street View machine-classified cell-year observations in mainland Jakarta, 2013–2019, weighted for sampling and imagery coverage. Each series gives the probability that the cell-year contains at least one classified adult of the indicated gender; it is not a share of person counts. The sample includes 2014, 2016 and 2017, which are excluded from the main opening regressions; 2014 has only 981 cell-years. Levels depend on classification and coverage. These are repeated observations of streets, not individuals; no confidence intervals are plotted.

already observed; it does not measure the probability of appearing on the street, the purpose of every outing or the constraints behind it.

Figure A6: Posture, company and activity among observed adults, by gender



Notes: Ratios of summed hand-labelled counts in 10,216 Jakarta panoramas, 2013–2024 excluding 2020–2021; no population adjustment. Posture and activity divide by all identified adults of the same gender, including non-motorised riders. Company divides by adults on foot or seated. Same-gender groups include adults with same-gender adults or only children; mixed-gender groups include adult women and men. The two trading-conditional rows in panel (a) divide by adults recorded as trading. With-children and street-vending indicators overlap other categories. Shares describe those visible under the labelling protocol, not residents or unique users. No uncertainty intervals are shown. The female share among seated adults in the text instead uses women plus men seated as its denominator.

A.10 Time Use, Safety and Care

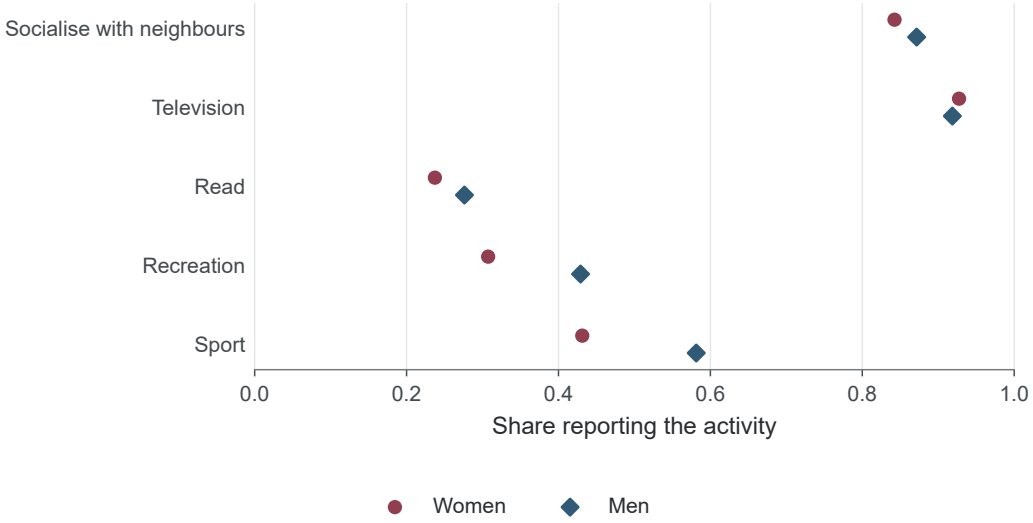
Statistics Indonesia's Happiness and Well-being Measurement Survey (SPTK) records time use, activities and household decisions in repeated cross-sections in 2014, 2017 and 2021. I restrict the data to Jakarta and apply the supplied survey weights. The public microdata identify the province but not finer locations, so they cannot be linked to the RPTRA spatial design. The comparisons provide context rather than estimates of opening effects.[17]

A.10.1 Reported Leisure and Activity Composition

Women report slightly more leisure than men in each wave. In 2021, the weighted means are 29.5 hours per week for women and 27.2 for men.[18] Approximately ninety per cent of both genders report watching television. Sport differs more: in 2021, 58 per cent of men and 43 per cent of women report it as a leisure activity. The reported gender gap ranges from 0.13 to 0.16 across waves and is significant at the 0.1 per cent level in each. Figure A7 shows the 2021 activity comparison.

The aggregate leisure comparison does not establish equal access to time outside the home. Employment differs by gender, and non-employed women report more committed time than non-employed men. Leisure may also be available at times or in places that make an outing difficult. These survey questions do not measure the visible non-work time in the accounting.

Figure A7: Share using leisure time for each activity, by gender, Jakarta 2021.



Notes: Statistics Indonesia's SPTK well-being survey, DKI Jakarta respondents in 2021. Points are survey-weighted shares (0–1) of respondents reporting each activity as a use of leisure time, calculated separately by gender among non-missing responses. Activities are not mutually exclusive; their shares need not sum to one. These are reported activities, not time shares or observations of activity outside the home. No confidence intervals are shown.

[17] The public files do not supply clustering information. The reported standard errors use probability weights without a released PSU or strata structure. The 2021 wave was collected during the pandemic; the earlier-wave comparisons are reported separately where available.

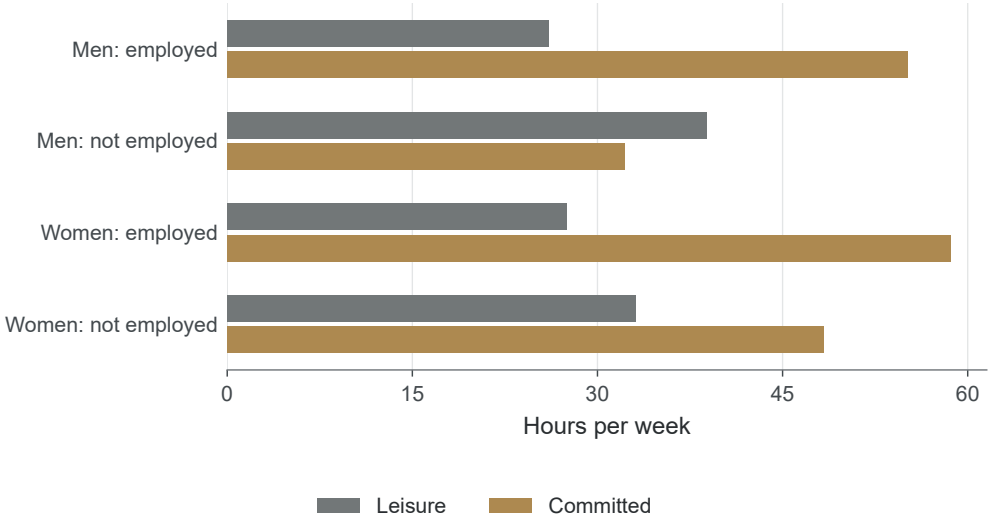
[18] The 2014 item adds leisure with and without family, unlike the single item in 2017 and 2021. Within-wave gender comparisons are reported, but levels are not treated as comparable across waves.

A.10.2 Employment and Committed Time

The leisure comparison changes after accounting for employment and demographics (Figure A8 and Table A6). The unconditional female-minus-male difference is 1.83 hours per week and is significant at five per cent; the adjusted difference is 0.05 hours and is not statistically significant. Employment is less common among women and is associated with less leisure. The adjustment therefore separates the raw comparison from one conditional on employment and the included characteristics, without identifying an employment effect.[19]

Committed time adds hours of work, housekeeping and schooling. Among non-employed respondents, women report about 48 committed hours per week and men 32, alongside leisure of 33 and 39 hours respectively. Among employed respondents, committed time is also higher for women, at 59 rather than 55 hours, while leisure is nearly the same. These are comparisons across respondents, not changes caused by entering or leaving employment. They describe time allocation that the street imagery does not directly observe.

Figure A8: Committed and leisure hours per week, by gender and employment status, Jakarta 2017 and 2021.



Notes: SPTK respondents in DKI Jakarta, pooled across 2017 and 2021 using supplied survey weights. Bars report mean hours per week by gender and employment status, using non-missing responses for each time measure. Committed time sums paid work, housekeeping and schooling; leisure is the reported leisure-hours item. Employment refers to the previous month in 2017 and previous week in 2021. The comparison is descriptive across respondents and is not the effect of employment. No confidence intervals are shown.

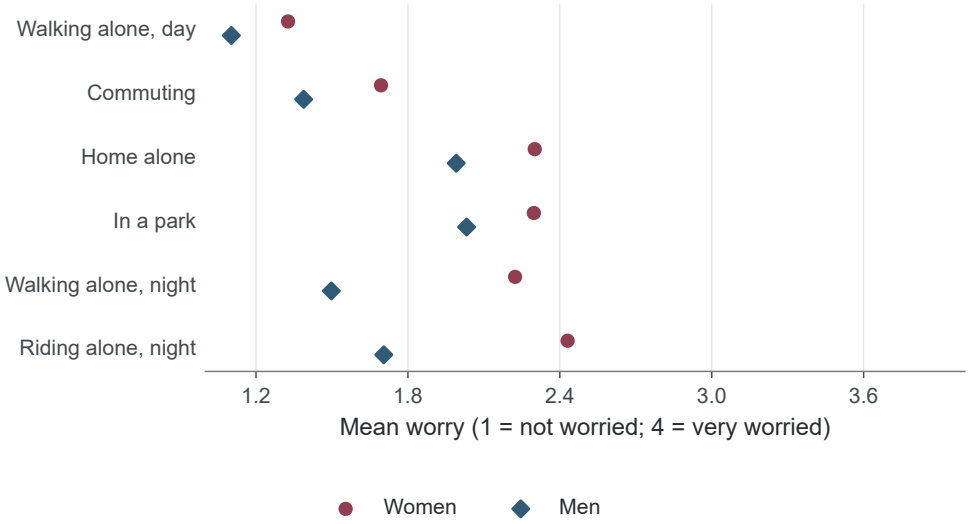
A.10.3 Reported Safety Concerns

Women report more worry than men in every setting measured in 2021, and the difference is larger after dark. On the four-point scale, moving from daytime to night-time walking alone raises mean worry by 0.89 points for women and 0.40 for men. Night-time walking and riding alone show the largest gender gaps: the means are 2.22 versus 1.50, and 2.43 versus 1.70, respectively (Figure A9).

[19] Employment refers to the previous month in 2017 and the previous week in 2014 and 2021. The regression table pools these waves and retains the differently worded 2014 leisure item.

These responses describe a possible influence on travel choices. They do not establish that safety explains the RPTRA response or reveal the gender gap in street occupancy at night.

Figure A9: Perceived safety, mean worry on a four-point scale, by setting and gender, Jakarta 2021.



Notes: SPTK respondents in DKI Jakarta, 2021. Points report survey-weighted mean worry by setting and gender among non-missing responses. The scale runs from 1 (not worried) to 4 (very worried); responses of not relevant are excluded. Settings are ordered by the gender difference. These are stated concerns, not observed victimisation or treatment effects; no confidence intervals are shown.

Table A6: Female leisure gap: attenuation by employment status

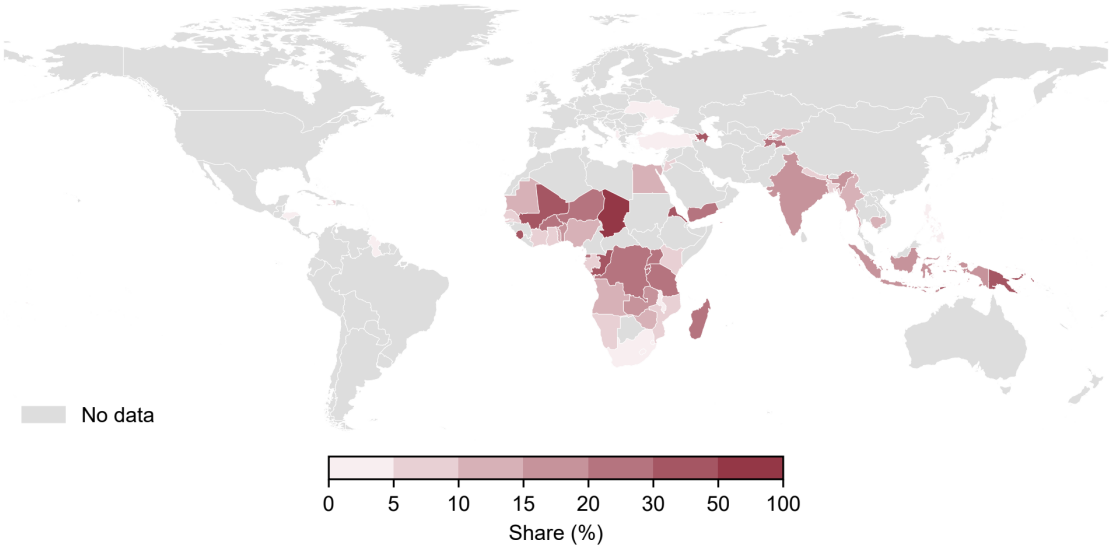
	(1) Raw	(2) + Employed	(3) + Controls
Female	1.830**	-1.644*	0.0519
	(0.818)	(0.894)	(0.937)
Employed		-8.424***	-6.460***
		(0.932)	(0.995)
Observations	3552	3552	3552

Notes: SPTK respondents in DKI Jakarta, pooled across 2014, 2017 and 2021. Survey-weighted linear regressions of weekly leisure hours include a female indicator and wave fixed effects in all columns. Column (2) adds employment; column (3) adds age, age squared, education, marital status and household income rank. Coefficients are hours per week. Employment refers to the previous month in 2017 and previous week in 2014/2021; the 2014 leisure item sums time with and without family and is not level-comparable across waves. Parentheses report linearised standard errors using probability weights only, without PSU or strata information. The adjustment is descriptive, not a causal employment effect. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

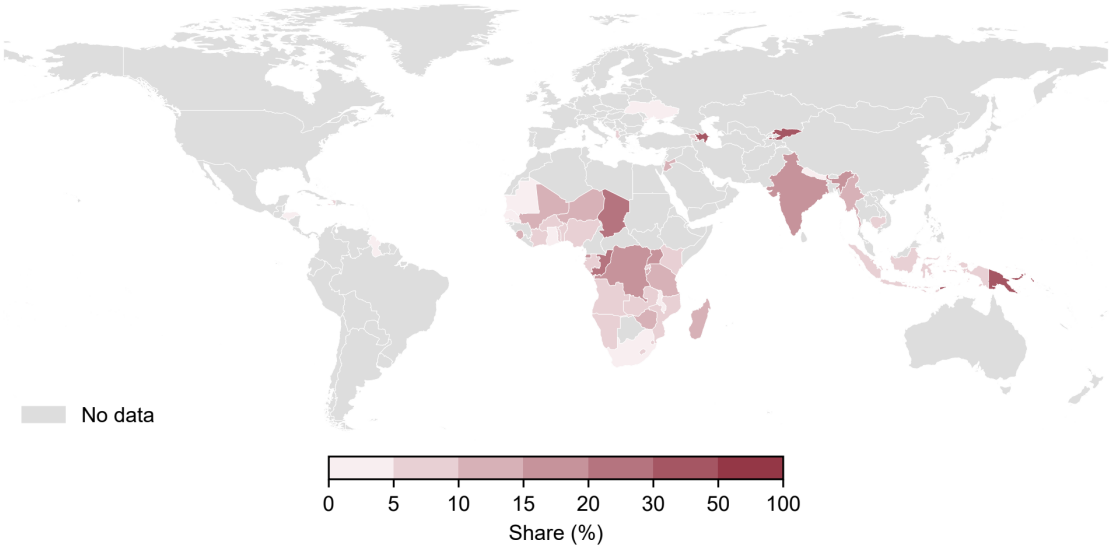
A.11 Supplementary Evidence on Gendered Constraints

The following figures provide cross-country context for Section 2.1. They cover attitudes towards a wife leaving without telling her husband (Figure A10), concern about violent crime (Figure A11) and female-to-male labour-force participation (Figure A12). Populations and reference years differ across indicators. The maps describe institutional conditions and reported attitudes, not the contributions of particular constraints to Jakarta's street gap.

Figure A10: Share reporting that beating is justified if a wife goes out without telling her husband.



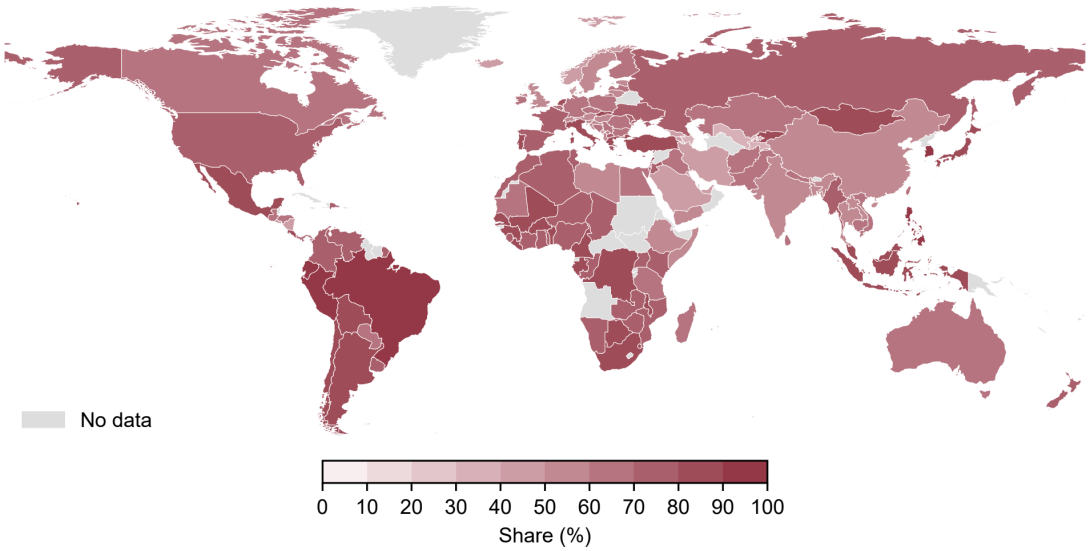
(a) Share of urban married women



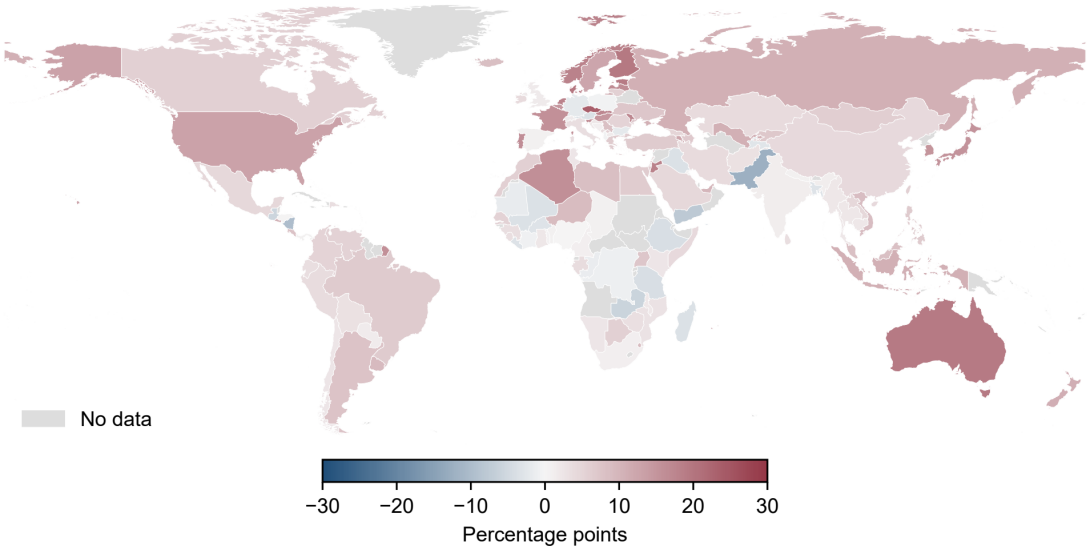
(b) Share of urban married men

Notes: Panels report percentages of urban married women and men who say that wife-beating is justified when a wife goes out without telling her husband. Source: the compiled Demographic and Health Surveys urban country extract. Countries retain their respective survey years; the maps are not a common-year cross-section. Missing countries are unshaded. Boundaries: Natural Earth 1:50 million.

Figure A11: Concern that violent crime could cause serious harm.



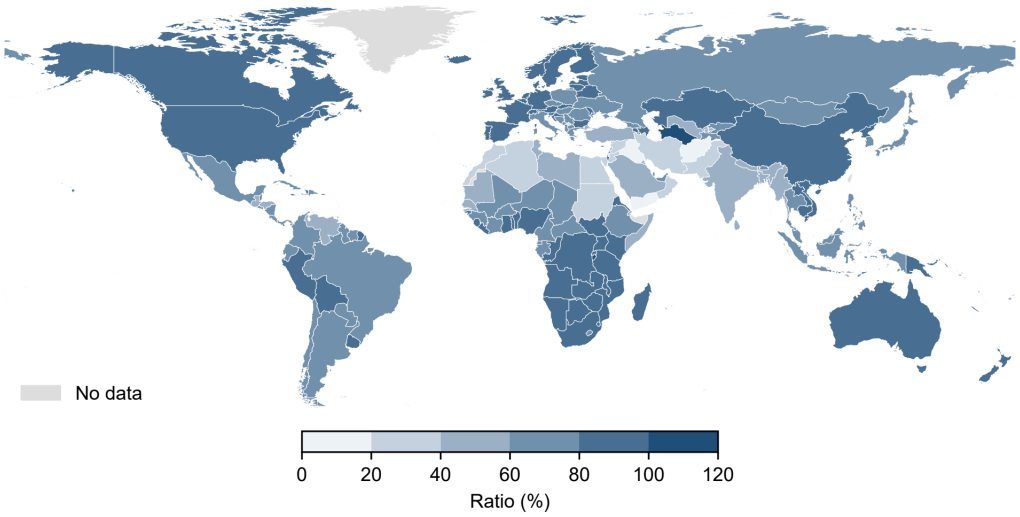
(a) Share of women reporting concern



(b) Gender gap in reported concern

Notes: Panel (a) reports the survey-weighted percentage of women expressing concern that violent crime could cause serious harm; panel (b) subtracts the corresponding male percentage, in percentage points. Source: World Risk Poll 2023, item WP20722, responses 1 or 2 among valid responses 1–3. Missing, refused and do-not-know answers are excluded separately by gender. The difference scale is centred on zero. Boundaries: Natural Earth 1:50 million.

Figure A12: Ratio of female to male labour force participation rate, per cent.



Notes: The map reports 100 times the female labour-force participation rate divided by the male rate. Source: World Bank extract of modelled ILO estimates, with 212 records dated 2024 and five dated 2021–2023. Country-specific years are retained; missing and unmatched areas are unshaded. Boundaries: Natural Earth 1:50 million.

B Travel Survey: Linkage, Samples and Purpose Definitions

B.1 Corrected Relational Links

I construct the survey sample from the SITRAMP household, person, trip, purpose and expansion-factor files. Complete household and person identifiers link the records; trip identifiers and reported diary totals provide additional checks. An earlier extract, `10Jabo2.dat`, attaches characteristics using sample and member numbers that repeat across zones. The audit finds that 97.8 per cent of its attached person records refer to another household. Even when trip fields match across deliveries, the attached individual characteristics in that extract are unreliable. I use corrected relational links for all gender-specific results.

The descriptive sample requires unique household, person, trip and weight identifiers; positive matched expansion factors; valid purposes; and agreement between linked trip counts and reported diary totals. Across the full roster, 154 trip rows have no person match and another 528 have conflicting linked identifiers. I exclude affected people, diary-total conflicts and household-linkage conflicts. These checks remove 5,862 people before the age and geographic restrictions. I retain zero-trip observations only when the diary is complete and valid.

The retained DKI sample aged 15–110 contains 93,730 men and 92,014 women. A sensitivity using the full observed roster gives non-work trip-making rates of 35.27 per cent for men and 48.23 per cent for women. The sensitivity includes unresolved records and serves as a check on the effect of these exclusions.

B.2 Purposes, Weights and Distinct Samples

The recorded purpose codes are work (1), school (2), shopping (3), business (4), private (5), return home (6) and other (7). For the descriptive comparisons, non-work comprises codes 3, 5 and 7, covering errands and private activities as well as leisure. Purpose-specific trip-making indicators can overlap because an adult may travel for several purposes on the same day. Per-adult counts retain validated non-travellers and count legs directed to a purpose, not complete outings. Person expansion factors weight trip-making and daily counts; trip expansion factors weight mean duration per leg.

The model uses a narrower set of trips: qualifying departures must start at home in the home neighbourhood and end within the 261 mainland neighbourhoods. Work includes business. Any-travel status is calculated from the complete diary before these restrictions. Known external destinations do not qualify for a model purpose. Missing or unmappable destinations are audited separately rather than automatically classified as non-travel.

Restoring two source-documented Pengadegan survey subzones retains 389 valid adults, giving the same roster of 185,744 adults across all 261 origins. When unresolved destination geography makes a person–purpose outcome ambiguous, that person is excluded from both estimation stages for the affected purpose. The main composite excludes 20 men and 21 women.

I distinguish households without recorded children from households whose child information is incomplete. The under-fifteen control has an explicit unknown category: 24,557 adults, or 13.2 per cent, belong to incomplete household rosters containing no observed child. A household roster beginning at age five cannot establish whether children below that age are present.

B.3 Travel Time and Completed Journeys

Clock values must be valid HHMM; 24:00 is the only permitted value with hour 24. An overnight adjustment is allowed only when departure is at or after 19:00 and arrival at or before 07:00. Durations must be positive and no longer than 300 minutes. Of 396,779 trips in the descriptive sample, 96 have invalid clocks and 114 exceed 300 minutes; 80 receive an overnight adjustment. Removing the duration cap changes purpose-by-gender means by less than half a minute.

I mark a person's total time for a category as missing when any contributing leg has an invalid duration, while validated non-travellers enter with zero minutes. Sample sizes consequently vary across time categories. Non-work daily totals do not allocate return-home time to an outbound purpose.

A completed journey starts at home and ends at the next recorded return. Reconstruction requires consecutive, unique trip numbers; valid place types; continuity of places and zones between adjacent legs; a first origin and final destination at home; and agreement between return-home purpose and home destination. An ambiguous sequence excludes the person's whole day.

The strict sequence sample retains 138,294 of 148,300 travellers, corresponding to 90.5 per cent of male and 96.6 per cent of female travellers. It contains 164,463 completed journeys. Of these, 139 include an invalid leg time, leaving 164,324 journeys for the duration comparison. Single-purpose journeys may visit several destinations of the same purpose; journeys with different non-home purposes are classified separately as mixed. Each journey receives its traveller's person expansion factor. Duration sums travel legs, including the return home, but excludes time spent at destinations.

Table B1: Completed home-to-home journey travel times, by purpose

Outcome	Men	Women	Difference [95% CI]	Adj. p	N_M	N_F
Work	90.2	90.9	0.8 [-0.9, 2.4]	0.362	37,644	15,694
School	82.2	83.4	1.2 [-0.5, 2.8]	0.328	10,184	9,555
Shopping	56.1	38.3	-17.8 [-19.4, -16.2]	< 0.001	4,802	25,008
Business	104.5	81.3	-23.2 [-29.6, -16.9]	< 0.001	2,808	879
Private	57.8	54.2	-3.6 [-5.1, -2.1]	< 0.001	13,781	14,161
Other	47.6	44.7	-2.9 [-5.2, -0.6]	0.038	3,181	4,307
Mixed purposes	137.3	115.3	-22.0 [-24.4, -19.7]	< 0.001	13,674	8,646
All journeys	89.0	67.3	-21.8 [-22.6, -20.9]	< 0.001	86,074	78,250

Notes: Corrected 2002 SITRAMP data for DKI adults aged 15–110. Travel minutes sum all legs of a completed home-to-home journey, including the return, and exclude destination time. Journeys receive the traveller's person expansion factor. A single-purpose journey may have several stops; a mixed journey has distinct non-home purposes. N_M and N_F count timed journeys, not unique travellers. Differences are women minus men. The 95 per cent normal intervals cluster by residential survey zone; Holm-adjusted p-values cover the eight rows. The strict sample contains 138,294 travellers and 164,463 journeys, of which 139 fail the timing rules. Destinations may lie outside the model's mainland system. Comparisons condition on valid sequences, with different retention rates by gender.

B.4 Trip Chaining

A chain visits at least two non-home destinations before returning home. Two separate home–destination–home journeys are not a chain, and a transport transfer within a recorded trip is not an activity stop. Under this definition, women chain less often in the corrected sequence sample. This differs from the higher frequency of complex female trip chains discussed by Ng and Acker (2018). My estimate conditions on completed sequences, with higher retention among female than male travellers, and is

specific to this definition of a chain. A place-code sensitivity relaxes agreement between home-purpose and home-place codes while retaining continuity checks.

Table B2 reports both the share of travellers making a chain and the share of completed journeys with multiple stops. Female chaining is lower under both denominators and in the work/business comparison; the relaxed place-code rule preserves that ordering. The chaining result concerns multiple stops within one home-to-home journey, not total trip frequency or travel distance. It cannot be used to infer differences in care responsibilities or scheduling constraints.

Table B2: Trip chaining under alternative denominators and sequence rules

Outcome	Men	Women	Difference [95% CI]	Adj. p	N_M	N_F
Travellers with a chained journey	23.6	16.1	-7.5 [-8.1, -6.8]	< 0.001	72,562	65,732
Journeys with multiple stops	20.0	13.6	-6.4 [-7.0, -5.9]	< 0.001	86,155	78,308
Work/business journeys with multiple stops	25.3	21.8	-3.5 [-4.4, -2.7]	< 0.001	51,513	20,607
Work travellers with a mixed journey	20.1	18.8	-1.3 [-2.1, -0.5]	0.009	48,944	20,046
Travellers with a chain: relaxed rule	23.1	16.2	-7.0 [-7.6, -6.4]	< 0.001	77,983	67,323

Notes: Corrected 2002 SITRAMP data for DKI adults aged 15–110. Male and female entries are percentages; differences are women minus men in percentage points. The 95 per cent normal intervals cluster by residential survey zone. Person expansion factors weight travellers or their journeys. N_M and N_F are unweighted denominator counts: travellers in rows 1, 4 and 5, and journeys in rows 2 and 3. A chain has at least two non-home stops; a mixed journey includes work/business and non-work destinations. Row 4 conditions on travellers with a work/business journey. Row 5 relaxes agreement between home-purpose and home-place codes but retains unique consecutive sequences, home-bounded days and continuity of places and zones. Holm adjustment uses the original family of 69 descriptive tests, not only these five rows. The clustering assumptions do not reproduce an official complex-survey design.

B.5 Inference and Spatial Reach

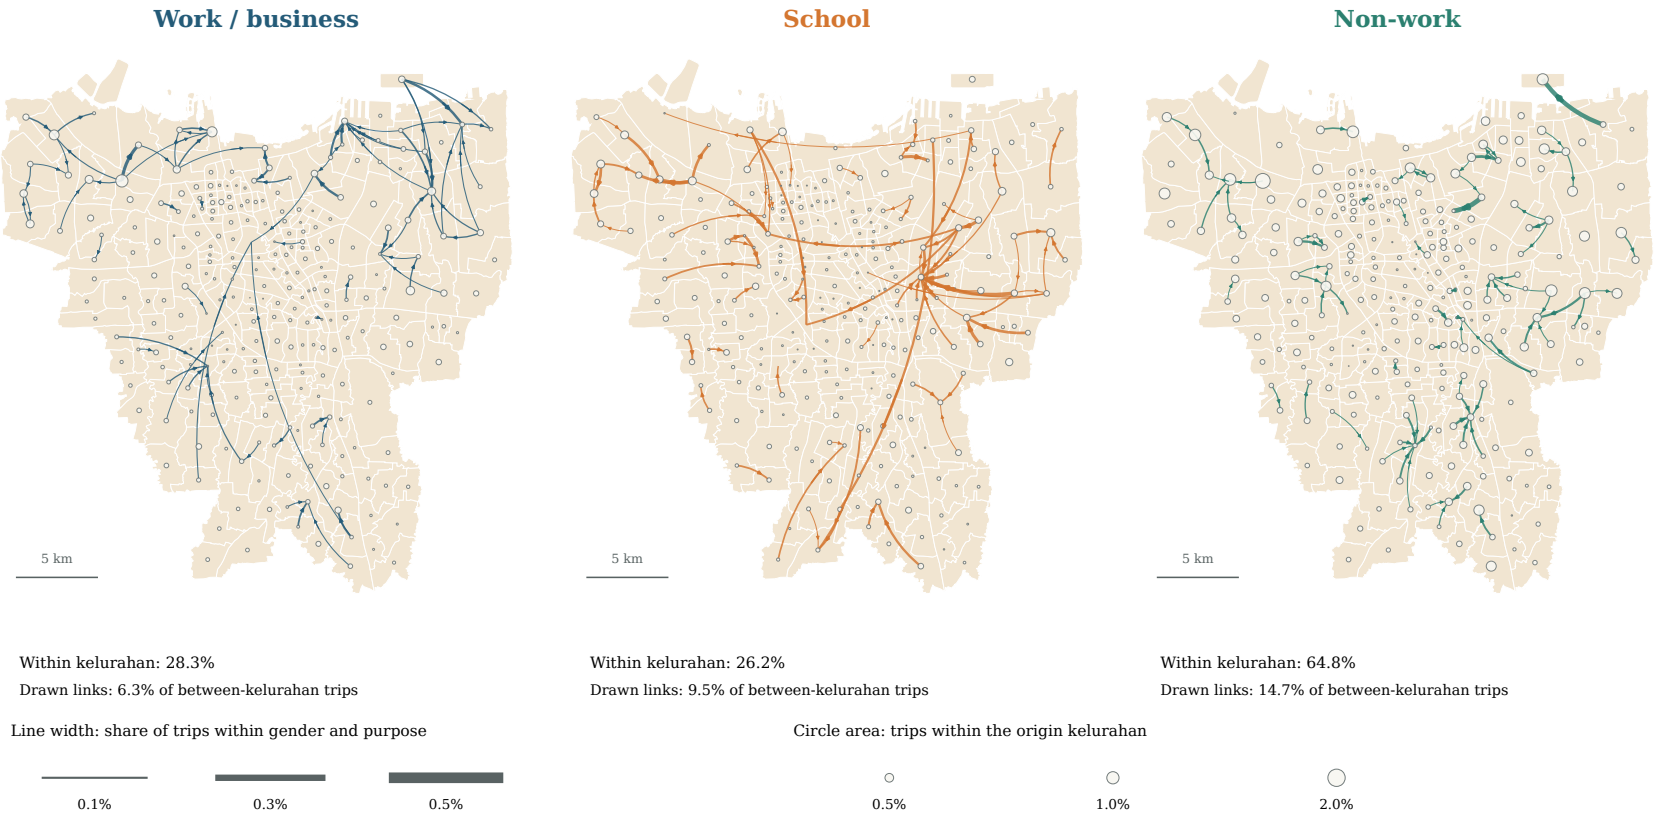
Descriptive tests cluster by residential survey zone, retaining women and men jointly and allowing repeated observations from a traveller. Household clustering is a sensitivity check. Variances use a finite-cluster sandwich correction and an asymptotic normal reference distribution. Holm adjustment applies to the stated family of tests; the journey-time family contains eight comparisons. The release does not document strata, PSU structure or replicate weights, so the reported errors are not full complex-survey design-based errors. Complete-record and sequence restrictions may also select which adults enter each comparison.

Neighbourhood-centroid distances proxy geographic separation, not routes. Internal distance is the mean separation of two points in an equal-area disc. Section 7.2 reports the purpose-specific decay estimates and their uncertainty; Appendix E.2 reports internal-distance sensitivities and destination-support checks. The estimates use the restricted home-departure sample rather than the complete-diary sample described above.

B.6 The Geography of Home-Originating Trips

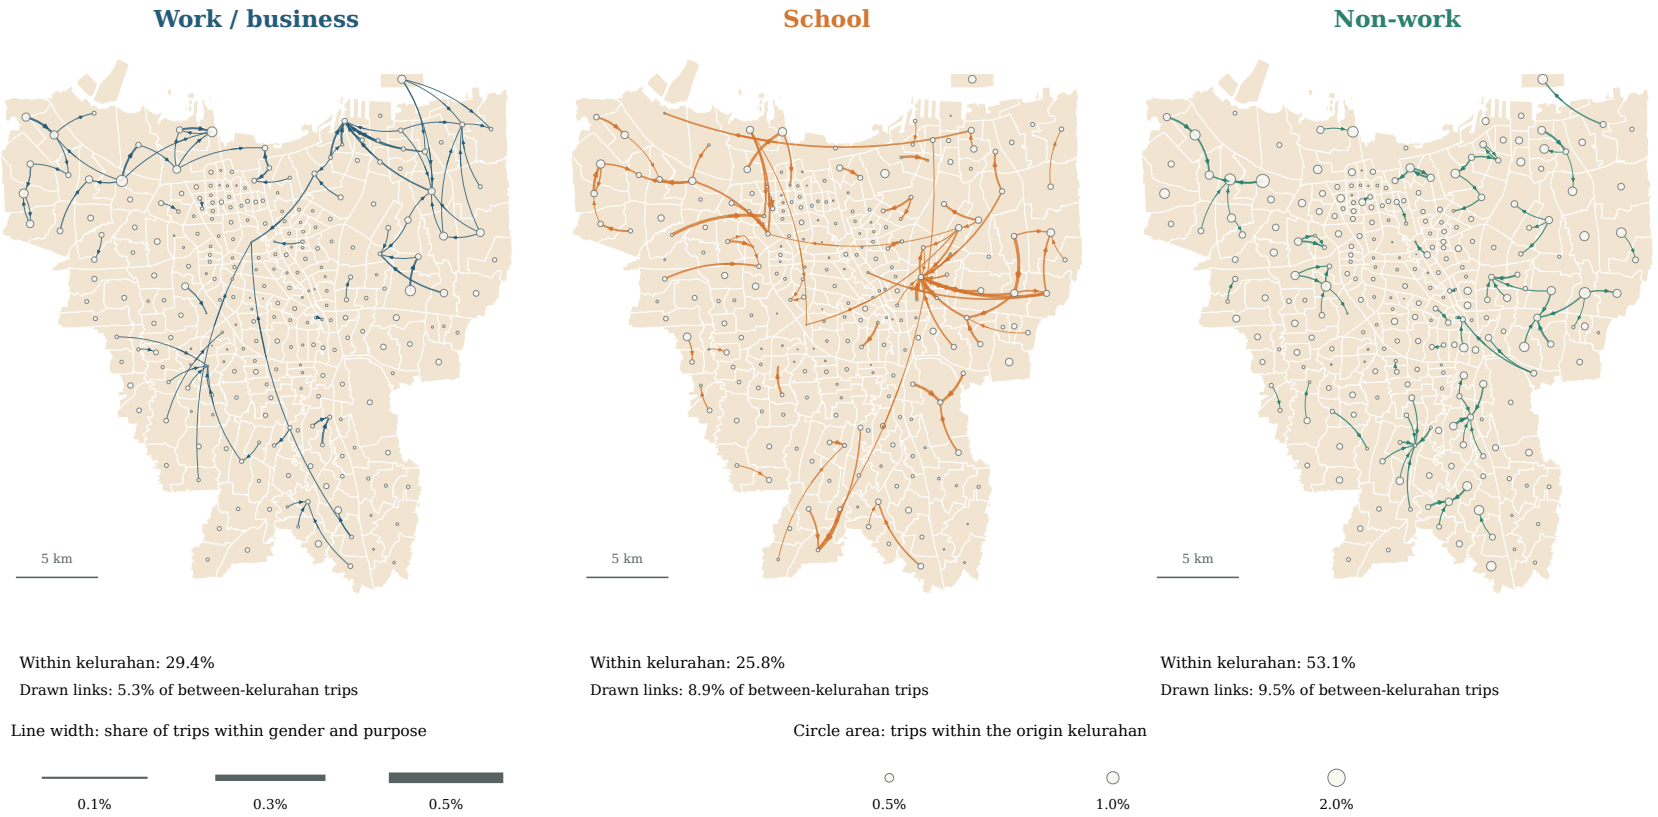
Figures B1 and B2 map the home-originating flows used in the spatial comparisons. The maps share a set of displayed connections and normalise flows within each gender and purpose. Women's non-work travel is concentrated more locally, while the work and school patterns are more similar across genders. These maps condition on recorded trips, rather than reporting the population probability of travelling.

Figure B1: Women's home-based trip flows, by purpose



Notes: 2002 SITRAMP, corrected person–trip linkage; mainland residents aged 15–110 with complete valid diaries. Trips depart from home in the home neighbourhood and end within the 261 mainland kelurahan. Return-home trips and known external destinations are excluded; unresolved potentially qualifying geography excludes the affected person–purpose. Work includes business; non-work comprises shopping, private and other purposes. Flows use trip expansion weights. Within each purpose, both gender maps use the union of each gender's 50 largest between-neighbourhood connections; origin and destination identifiers break ties. Positive flows on these connections are drawn; omitted connections need not be zero. All within-neighbourhood flows appear as circles. Line widths and circle areas show shares of total weighted trips within gender and purpose, on common scales. Panel footers report within-neighbourhood shares and coverage of between-neighbourhood volume. Arrows connect centroids (interior points where necessary), not routes. Boundaries follow the current model crosswalk, not reconstructed 2002 geography. These descriptive shares have no uncertainty intervals and do not measure population trip-making rates.

Figure B2: Men's home-based trip flows, by purpose



Notes: 2002 SITRAMP, corrected person–trip linkage; mainland residents aged 15–110 with complete valid diaries. Trips depart from home in the home neighbourhood and end within the 261 mainland kelurahan. Return-home trips and known external destinations are excluded; unresolved potentially qualifying geography excludes the affected person–purpose. Work includes business; non-work comprises shopping, private and other purposes. Flows use trip expansion weights. Within each purpose, both gender maps use the union of each gender's 50 largest between-neighbourhood connections; origin and destination identifiers break ties. Positive flows on these connections are drawn; omitted connections need not be zero. All within-neighbourhood flows appear as circles. Line widths and circle areas show shares of total weighted trips within gender and purpose, on common scales. Panel footers report within-neighbourhood shares and coverage of between-neighbourhood volume. Arrows connect centroids (interior points where necessary), not routes. Boundaries follow the current model crosswalk, not reconstructed 2002 geography. These descriptive shares have no uncertainty intervals and do not measure population trip-making rates.

C Accounting for the Gender Gap: Derivations and Measurement

Starting from the work identity in (1), I derive the conditional minimum residual in Section 4.2, describe the symmetric decomposition and document its inputs. I then vary the visibility and commuting assumptions and the treatment of unidentified adults, and repeat the accounting for the post-COVID window.

C.1 Derivation of the Conditional Minimum Residual

Let S_f be the female share of identified visible adults and N_g the adult population of gender g. The population adjustment gives the per-adult presence ratio

$$R \equiv V_m/V_f = (N_f/N_m)(1 - S_f)/S_f. \tag{C1}$$

$$\underline{\ell} = (R - \rho_W)/(R - 1), \qquad 1 < \rho_W < R. \tag{C2}$$

The positive minimum in (C2) applies when measured work exposure lies between parity and the observed occupancy ratio.

Write $W_g = p_g \bar{v}_g$ for measured work exposure, so that $V_g = W_g + \ell_g$. The two ratios are

$$R \equiv \frac{V_m}{V_f}, \qquad \rho_W \equiv \frac{W_m}{W_f}.$$

I construct R from the hand-labelled imagery and population adjustment, and ρ_W from the labour-force and commuter inputs under the stated visibility assignments. For fixed inputs, the residual share of the gap is

$$L(\ell_f) \;\equiv\; \frac{\ell_m - \ell_f}{V_m - V_f}. \tag{C3}$$

Proposition 1 (Conditional minimum residual)**.** *Suppose $\ell_f \geq 0$, $W_f > 0$ and $1 < \rho_W < R$. Then $L(\ell_f)$ is increasing in ℓ_f, and*

$$L(\ell_f) \;\geq\; L(0) \;=\; \frac{R - \rho_W}{R - 1} \;>\; 0.$$

Proof. Hold W_f, R and ρ_W fixed and vary ℓ_f. Maintaining $V_m = RV_f$ determines male residual time:

$$\ell_m \;=\; R\,(W_f + \ell_f) - W_m \;=\; W_f(R - \rho_W) + R\,\ell_f.$$

The total gap is $V_m - V_f = (R - 1)(W_f + \ell_f)$. Substitution into (C3) gives

$$L(\ell_f) \;=\; \frac{W_f(R - \rho_W) + (R - 1)\,\ell_f}{(R - 1)(W_f + \ell_f)}. \tag{C4}$$

Differentiating with respect to ℓ_f,

$$L'(\ell_f) \;=\; \frac{W_f\,[(R - 1) - (R - \rho_W)]}{(R - 1)(W_f + \ell_f)^2} \;=\; \frac{W_f(\rho_W - 1)}{(R - 1)(W_f + \ell_f)^2} \;>\; 0$$

under the stated condition $1 < \rho_W < R$. The minimum is therefore at $\ell_f = 0$, where (C4) equals

$(R - \rho_W)/(R-1) > 0$. As ℓ_f grows without bound, the residual share approaches one. □

The conditional minimum residual depends only on the two ratios and is unchanged by a common rescaling of visible minutes. It minimises the residual share over female visible time outside measured work, holding work inputs fixed. Positive female residual time requires greater male residual time to reproduce R, increasing the share.

When $\rho_W > R > 1$, setting $\ell_f = 0$ would imply negative male residual time. Requiring both residual components to be non-negative instead gives the signed minimum $(R - \rho_W)/[\rho_W(R-1)]$, attained at $\ell_m = 0$. There is no positive minimum in that case.

The baseline 2015–2019 values are approximately $R = 2.81$ and $\rho_W = 2.26$, giving the reported 30.2 per cent minimum using unrounded inputs. The 2022–2024 values are approximately $R = 3.36$ and $\rho_W = 2.45$, giving 38.7 per cent. Tables C3 and C2 report the inputs. Across the two windows, the observed occupancy ratio rises more than the measured-work ratio under the respective input assignments.

C.2 Accounting Counterfactuals and Positive Non-work Exposure

The main-text employment counterfactual uses the 2015–2019 presence ratio and 2018-based work inputs. Let $k = p_m/p_f$. Raising women's employment rate to p_m while retaining their visible exposure per employed person gives $W_f^{\text{eq}} = kW_f$. The work-exposure ratio falls to $\rho_W/k = 1.409$. The work-exposure comparison does not depend on the amount of non-work visible time.

A prediction for total occupancy also requires the residual components. If both remain at their baseline values and $x = \ell_f/W_f$, then

$$R^{\text{eq}}(x) = \frac{W_m + \ell_m}{kW_f + \ell_f} = \frac{R(1+x)}{k+x}. \tag{C5}$$

For $k > 1$, $R^{\text{eq}}(x)$ increases with x. It is 1.75 at $x = 0$ and approaches the baseline 2.81 as x grows. The predicted total-occupancy ratio thus depends on x, even with residual activity held fixed. Employment-induced changes in other activities are outside this calculation.

Table C1 considers a different comparison, retaining baseline employment and varying assumed female non-work exposure. For $x > 0$, the identity $R = (W_m + \ell_m)/(W_f + \ell_f)$ implies

$$\frac{\ell_m}{\ell_f} = R + \frac{R - \rho_W}{x}, \qquad L(x) = \frac{R - \rho_W + (R-1)x}{(R-1)(1+x)}. \tag{C6}$$

When $R > \rho_W$, the residual-exposure ratio decreases with x, while its share of the total gap increases when $\rho_W > 1$. The residual can therefore account for more of the total gap even as its own male-to-female ratio declines. At $x = 0$, ℓ_m/ℓ_f is undefined but $L(0)$ is defined. The scenarios hold the observed ratios fixed and are not estimated responses to a change in employment.

C.3 The Symmetric Decomposition

Employment and exposure per employed worker enter multiplicatively. I allocate their joint contribution symmetrically across the two gender orderings. Decomposing first employment and then exposure gives a different allocation of their interaction than reversing the order. Averaging those two orderings

Table C1: Non-work exposure and the accounting residual: conditional scenarios

Female non-work exposure $x = \ell_f / W_f$	Non-work ratio ℓ_m / ℓ_f	Residual share of the gap (%)
0 (conditional minimum)	–	30.2
0.25	4.99	44.2
0.50	3.90	53.5
1.00	3.35	65.1

Notes: Conditional calculations use 2015–2019 hand-labelled pedestrian counts, SAKERNAS 2018 employment and hours, and 2023 work-location and walking-duration inputs for mainland Jakarta adults aged 15 and over. $R = 2.807$ is male-to-female visible time per adult; $\rho_W = 2.261$ is the measured-work exposure ratio. Each row varies $x = \ell_f / W_f$ while holding both ratios fixed. The residual-exposure ratio divides male by female visible time outside measured work and is undefined at $x = 0$. The residual share is $100[R - \rho_W + (R-1)x]/[(R-1)(1+x)]$. Calculations retain the middle visibility assignment, common within-category gender visibility, exclusion of unidentified adults and walk-only commuting. These are assumed scenarios, without sampling intervals.

gives

$$V_m - V_f = \underbrace{(p_m - p_f)\,\bar{\bar{v}}}_{\text{employment}} + \underbrace{\bar{p}\,(a_m - a_f)}_{\text{occupational visibility and hours}} + \underbrace{\bar{p}\,(c_m - c_f)}_{\text{commuting}} + \underbrace{(\ell_m - \ell_f)}_{\text{non-work}}, \tag{C7}$$

where $\bar{p} = (p_m + p_f)/2$, $\bar{\bar{v}} = (\bar{v}_m + \bar{v}_f)/2$ and $\bar{v}_g = a_g + c_g$. Here a_g denotes visible working minutes per employed person. The first ordering is the one used by Kleven and Landais (2017); I average the two orderings rather than selecting either gender as the sole reference. The resulting contributions are symmetric in the two orderings, although other groupings of the inputs could yield different allocations.

Visible working minutes are written as $a_g = h_g \phi_g$. Applying the same symmetric rule to hours and visibility gives

$$a_m - a_f = \underbrace{(h_m - h_f)\,\bar{\phi}}_{\text{hours}} + \underbrace{\bar{h}\,(\phi_m - \phi_f)}_{\text{visibility}}. \tag{C8}$$

Figure 5 divides the resulting employment, hours, visibility and commuting terms by the total gap, evaluated at $\ell_f = 0$. Contributions retain their signs. A margin with greater female than male exposure offsets the gap; commuting has that sign in the baseline.

C.4 Constructing Each Component

A common counting rule. I construct work exposure to match the pedestrian-only imagery outcome. It counts people on foot or seated outside a vehicle and excludes people travelling in or on vehicles, including non-motorised ones. The hand-labelling protocol in Section A.1 records non-motorised riders separately, allowing that exclusion. The visibility assignments must follow the same rule, including the treatment of work performed beside rather than inside a vehicle. For plant and machine operators in the combined street/roadside/itinerant category, the low, middle and high visible fractions are 0, 0.25 and 0.50. These fractions allow some off-vehicle work or waiting; they are assumptions, not measured waiting durations.

Employment, p_g. Employment is the SAKERNAS employment-to-population ratio for adults aged 15 and over, with no upper age limit, in August 2018 and August 2023. This is the survey age convention used to approximate the adult category in the imagery. Table A5 reports headline and main-activity employment so that the two definitions can be compared.

Working minutes, h_g. Working time comes from usual weekly hours in the main job, capped at 84 hours. The implementation multiplies weekly hours by $60/7$ to obtain minutes per calendar day. There are no missing usual-hour values among employed adults in the retained 2018 and 2023 cleaned samples.

Visibility of work, ϕ_g. The 2023 work-location variable distinguishes five groups: the worker's own home, an employer's or client's home, an office or factory, a mall or covered market, and a street, roadside or itinerant location. Each group receives an assigned visible fraction, with occupation used where the location category is too broad.

Work inside a vehicle receives low visibility. Without this distinction, driving would be treated as highly visible simply because it takes place outdoors. I therefore separate machine and vehicle operators from other street workers. Street-front work receives high visibility. Where setting determines exposure, it takes precedence over occupation: a seller in an open market is treated differently from a salesperson inside a mall.

The fractions are not observed, so the accounting is repeated under low, middle and high assignments (Figure C1). Across the specified assignments, higher work exposure accounts for more of the gap and leaves a smaller residual. Other gender-specific visibility changes need not have the same effect. The low/middle/high fractions are $0/0.15/0.30$ for malls or markets in buildings and for other locations; 1 for markets without buildings and for non-operator street/roadside/itinerant work; $0/0.25/0.50$ for street/roadside/itinerant plant and machine operators; and 0 for the remaining locations. For 2018, the code applies the 2023 survey-weighted fraction by gender and occupation, falling back to the pooled occupation fraction when necessary. It then multiplies each 2018 respondent's hours by this transported fraction.

Let θ_g^k be the share of total employed working hours in category k and α_g^k its assigned visible fraction. For categories with a common assignment, the equivalent hours-weighted fraction is $\phi_g = \sum_k \theta_g^k \alpha_g^k$. The gender difference separates as

$$\phi_m - \phi_f \;=\; \underbrace{\textstyle\sum_k (\theta_m^k - \theta_f^k)\,\bar{\alpha}^k}_{\text{composition}} \;+\; \underbrace{\textstyle\sum_k \bar{\theta}^k (\alpha_m^k - \alpha_f^k)}_{\text{within category}}. \tag{C9}$$

Bars denote gender averages within each category. The baseline assignment varies by category but not gender, imposing $\alpha_m^k = \alpha_f^k$. Its within-category term is therefore zero, and the visibility contribution comes entirely from employment composition across categories. Allowing lower female visibility within categories reduces W_f, raises ρ_W and lowers the conditional minimum residual with the other inputs fixed. The implementation defines $a_g = (60/7)\mathbb{E}_w[H_i\alpha_i \mid g, \text{employed}]$, $h_g = (60/7)\mathbb{E}_w[H_i \mid g, \text{employed}]$ and $\phi_g = a_g/h_g$, where H_i is capped usual weekly hours and w the survey weight. In 2018, α_i is the transported gender-by-occupation fraction, so its within-occupation gender difference is not constrained to zero even though the underlying 2023 location assignments are common.

Visible commuting minutes, c_g. The surveys do not separately record walking and waiting within mixed-mode journeys. The walking-only benchmark averages assigned daily walking exposure over employed adults:

$$c_g = \mathbb{E}_w[2t_i(d_i/7)\,\mathbf{1}\{\text{daily commuter and walk-only}\} \mid g, \text{employed}]\,, \tag{C10}$$

where t_i is calibrated one-way walking duration and d_i the reported number of working days per week. Non-commuters and non-walkers contribute zero. Duration uses gender-specific 2023 commuter calibrations: gender-by-duration-band values for 2018 and gender-specific walking means for 2023.

Walking-only exposure omits access and waiting time on mixed-mode journeys. A sensitivity adds an allowance ω to the street-exposed transit share:

$$c_g(\omega) = c_g(0) + \Omega_g \, s_g^{\text{transit}} \, r_g^{\text{commute}} \, d_g^{\text{transit}} / 7. \tag{C11}$$

Here Ω_g is the calibrated daily out-of-vehicle allowance, s_g^{transit} the transit share among commuters, r_g^{commute} the commuting rate among employed adults and d_g^{transit} transit commuting days per week. The reported female street-exposed transit share is approximately three and a half times the male share (Table A5, panel C). Under the common allowance and the other baseline inputs, the reported sensitivity raises female relative to male exposure, lowers ρ_W and raises the conditional minimum residual. Transfers and roadside waiting for app-hailed motorcycles are additional possible sources of visible time, but their contributions are not separately measured. The baseline sets $\omega = 0$. Adding the specified transit allowance raises the conditional minimum residual, although other commuting conventions could change it differently.

Figure C3 reports the operational comparison between no correction, a calibrated out-of-vehicle allowance and its upper calibration limit. The caption describes the duration-difference calibration and commuting-rate scaling. They vary the commuting inputs rather than estimating a sampling distribution for the residual.

Non-work visible minutes, ℓ_g. No source used here directly measures both the duration and street visibility of non-work activity over the imagery window. The snapshots also do not follow individuals after they leave view. I therefore leave non-work visible minutes as a residual, with the implied minimum depending on the work inputs and their comparability with the imagery.

C.5 Adults of Unidentifiable Gender

Between eight and ten per cent of hand-labelled adults have unidentifiable gender (Table A5, panel E). Since R uses identified adults, assigning this group differently changes the conditional minimum residual. Figure C1 retains a positive minimum when these adults are excluded or all assigned to men under the reported visibility assignments. Assigning all to women removes the positive minimum under the middle and high assignments.

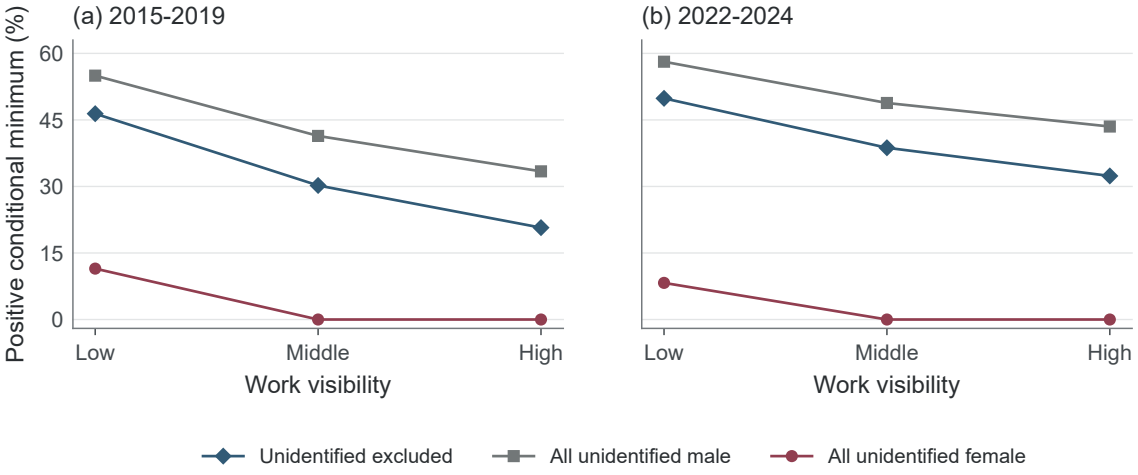
Exclusion defines the female share among identified adults. Extending it to all adults requires an assumption about non-identification. Distance, truncation and viewing angle may affect whether gender can be assigned, but do not reveal the unidentified group's gender composition without independent labels. The all-female and all-male assignments are sensitivity cases, not estimates of the unidentified population. I impose no direction for the bias from excluding it.

C.6 Sensitivity

Figure C1 combines alternative work-visibility assignments with alternative treatments of unidentified adults. The remaining exercises vary female within-category visibility and commuting exposure.

Each changes one part of the measurement while holding the other baseline inputs fixed.

Figure C1: Conditional minimum residual, by visibility assignment and treatment of unidentified adults



Notes: Hand-labelled pedestrian imagery in 2015–2019 and 2022–2024 is paired with survey-weighted SAKERNAS 2018 and 2023 adult inputs respectively. Work-location and walking-duration information comes from 2023. The calculations vary low, middle and high category visibility fractions, common across genders, and the assignment of unidentified adults: excluded, all male or all female. Commuting is walk-only. For $1 < \rho_W < R$, the minimum is $100(R - \rho_W)/(R - 1)$ at zero female residual time, where R is the population-adjusted presence ratio and ρ_W the work-exposure ratio. Cases without a positive residual minimum are displayed at zero as a no-positive-bound indicator, not as an estimated zero residual. The feasibility restriction outside this range is described in Section C.1. These are deterministic sensitivities without sampling intervals.

Work visibility. Low, middle and high assignments allow for different visible fractions within the broad workplace categories. Across these scenarios, the conditional minimum residual falls as assigned work visibility rises and remains positive when unidentified adults are excluded. Workers within a category may still differ: a roadside vendor and an itinerant courier need not spend the same fraction of working time in view.

Female within-category visibility. To relax the common-visibility restriction, let $\lambda \in (0, 1]$ scale female visible working time while leaving commuting and all male inputs fixed:

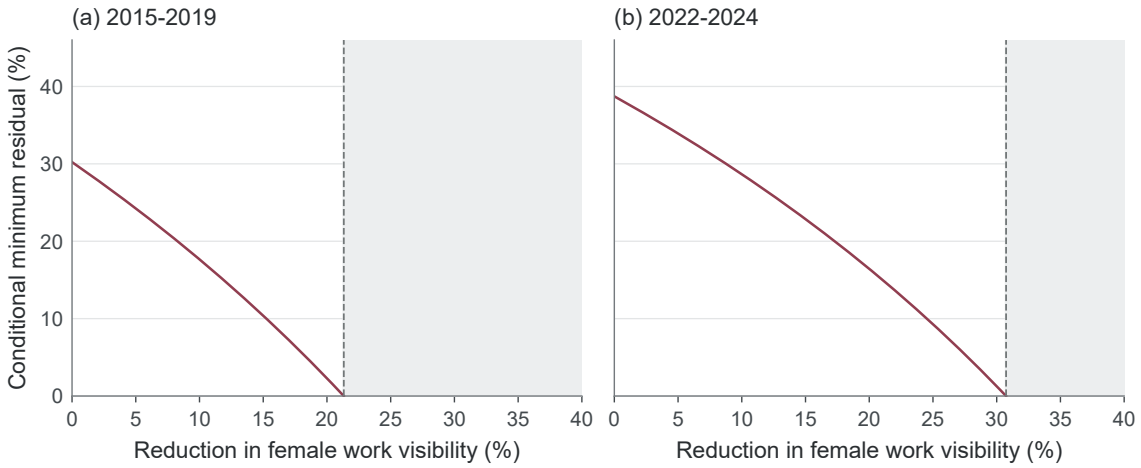
$$W_f(\lambda) = p_f(\lambda a_f + c_f), \qquad \rho_W(\lambda) = \frac{W_m}{W_f(\lambda)}.$$

Figure C2 reports the resulting conditional minimum residual. It reaches zero after a reduction of 21.3 per cent in assigned female visible working time in 2015–2019, and 30.7 per cent in 2022–2024. The thresholds quantify sensitivity to the visibility assignment; they are not estimates of its measurement error.

Unidentified adults. The alternative assignments in Section C.5 change the presence ratio rather than work exposure. Assigning all unidentified adults to men increases the conditional minimum residual; assigning all to women lowers it and can remove a positive minimum. The baseline exclusion rule is therefore part of the estimand, not evidence that unidentified adults share the identified group's composition.

Out-of-vehicle commuting time. Figure C3 adds an allowance for walking and waiting within street-exposed transit journeys. Under that calibration, female relative exposure rises and the conditional

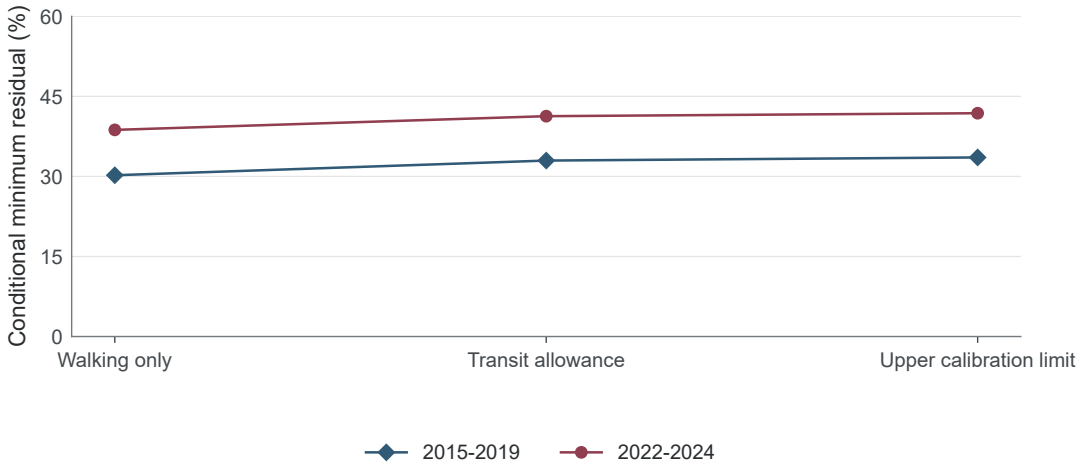
Figure C2: Sensitivity of the accounting residual to female work visibility



Notes: Hand-labelled Jakarta pedestrian imagery in 2015–2019 and 2022–2024 is paired with survey-weighted SAKERNAS 2018 and 2023 adult inputs, respectively. Work-location and walking-duration information comes from 2023 surveys. Each panel scales assigned female visible working time by λ, holding employment, hours, female commuting and all male inputs fixed. The horizontal axis is $100(1-\lambda)$; the vertical axis is $100[R-\rho_W(\lambda)]/(R-1)$ at zero female non-work visible time, where R is population-adjusted male/female presence and ρ_W male/female work exposure. Inputs retain mid visibility, exclusion of unidentified adults and walk-only commuting. Dashed lines mark the 21.3 and 30.7 per cent reductions at which $\rho_W = R$. Negative values are omitted, not set to zero; shaded regions have no positive lower bound. These are deterministic sensitivities, not estimates of visibility error or confidence intervals.

minimum residual increases. The walk-only baseline gives the smallest residual within this particular comparison; it need not be conservative under other measurement changes.

Figure C3: Conditional minimum residual under an allowance for out-of-vehicle commuting time



Notes: Hand-labelled Jakarta pedestrian imagery in 2015–2019 and 2022–2024 is paired with survey-weighted SAKERNAS 2018 and 2023 adult inputs, respectively. Work-location and walking-duration information comes from 2023 surveys. The vertical axis is the conditional minimum non-work share of the male-minus-female presence gap, in per cent, retaining mid work visibility and excluding unidentified adults. No correction counts walk-only commuting. Upper bound adds an allowance for out-of-vehicle time on street-exposed transit, based on the duration difference from private motor-vehicle journeys of the same reported distance, excluding routine intermediate stops and scaled by daily commuting rates. The difference bounds out-of-vehicle time only under a nonnegative transit speed penalty. Upper 95 per cent limit uses the upper calibration limit, not a confidence interval for the accounting residual. Ojek, ride-hailing, taxi, shuttle and becak are excluded from the transit allowance. Female non-work visible time is set to zero.

C.7 Post-COVID Accounting

The post-COVID accounting pairs SAKERNAS 2023 with hand-labelled imagery from 2022–2024. It retains the same identity, middle visibility assignment, exclusion of unidentified adults and walk-only commuting convention. Table C2 gives the inputs and Figure C4 separates the four work contributions and the residual.

The population-adjusted presence ratio is 3.364 and the measured-work ratio is 2.448. At zero female residual time, the conditional minimum residual is 38.7 per cent, compared with 30.2 per cent in the earlier window. A 30.7 per cent reduction in assigned female visible working time removes the positive post-COVID minimum, holding the other inputs fixed (Figure C2). The threshold is a sensitivity calculation, while the comparison across windows combines changes in both survey and imagery inputs rather than isolating an effect of the pandemic.

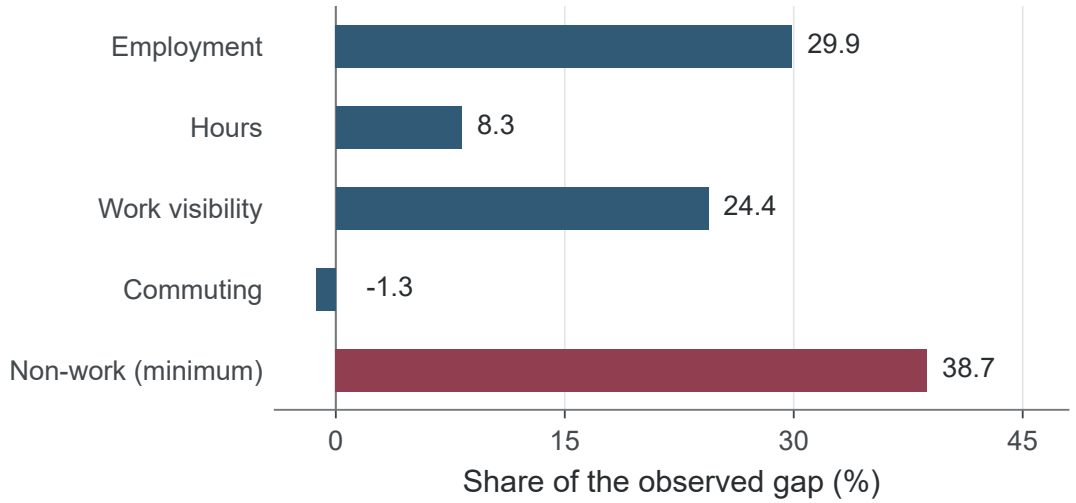
Table C2: Components of the street-presence identity, post-COVID accounting

Imagery window	SAKERNAS 2023 2022–2024
Employment rate, men	0.740
Employment rate, women	0.478
Visible work minutes per employed man	54.250
Visible work minutes per employed woman	32.205
Visible commute minutes per employed man	3.316
Visible commute minutes per employed woman	4.171
Visible work time per adult man	42.596
Visible work time per adult woman	17.397
Visible work-time ratio	2.448
Observed presence ratio	3.364
Share of observed gap (%), at zero female non-work time	
Employment	29.9
Work visibility and hours	32.7
Commuting	-1.3
Non-work presence (minimum)	38.7

Notes: SAKERNAS August 2023, adults aged 15 and over in mainland DKI, survey weighted. The presence ratio uses 2022–2024 hand-labelled imagery and adjusts by the adult population gender ratio. Visible work minutes use the mid work-location assignment. Commuting counts walk-only journeys with the duration calibration described in Section C.4. Work visibility and hours are combined in the table and separated in the figure. The non-work share is evaluated at $\ell_f = 0$.

For comparison, Table C3 reports the 2018-based inputs used in the main text.

Figure C4: Contributions to the street-presence gap, post-COVID accounting



Notes: Symmetric decomposition using SAKERNAS 2023 and 2022–2024 hand-labelled imagery, evaluated at zero female non-work visible time. Bars sum to 100 per cent, including the negative commuting contribution. Construction follows Section C.3. The residual minimum is conditional on the mid-visibility assignment, common within-category visibility across genders, exclusion of unidentified adults and the baseline commuting convention.

Table C3: Components of the street-occupancy identity, 2018-based accounting

Imagery window	SAKERNAS 2018 2015–2019
Employment rate, men	0.727
Employment rate, women	0.453
Visible work minutes per employed man	53.861
Visible work minutes per employed woman	36.810
Visible commute minutes per employed man	2.977
Visible commute minutes per employed woman	3.528
Visible work time per adult man	41.338
Visible work time per adult woman	18.285
Visible work-time ratio	2.261
Observed presence ratio	2.807
Share of observed gap (%), at zero female non-work time	
Employment	40.3
Work visibility and hours	30.5
Commuting	-1.0
Non-work presence (minimum)	30.2

Notes: Employment rates and hours are from SAKERNAS August 2018, adults 15 and over, mainland DKI, weighted. The work-visibility assignment uses the 2023 work-location information at the baseline; walk durations are calibrated using the 2023 commuter survey. The imagery window is 2015–2019. Visible working minutes are usual hours multiplied by a visibility fraction assigned to the reported work location and, for street work, to occupation; the table shows the middle visibility assignment. Visible commuting minutes count walk-only journeys. R is the male-to-female ratio of visible presence per adult from the hand-labelled imagery, counting adults on foot or seated, excluding those on non-motorised vehicles and those whose gender could not be identified, adjusted by the adult gender ratio. The minimum residual share is $(R - \rho_W)/(R - 1)$ from (C2). The residual minimum is conditional on the mid-visibility assignment, common within-category visibility across genders, exclusion of unidentified adults and the baseline commuting convention.

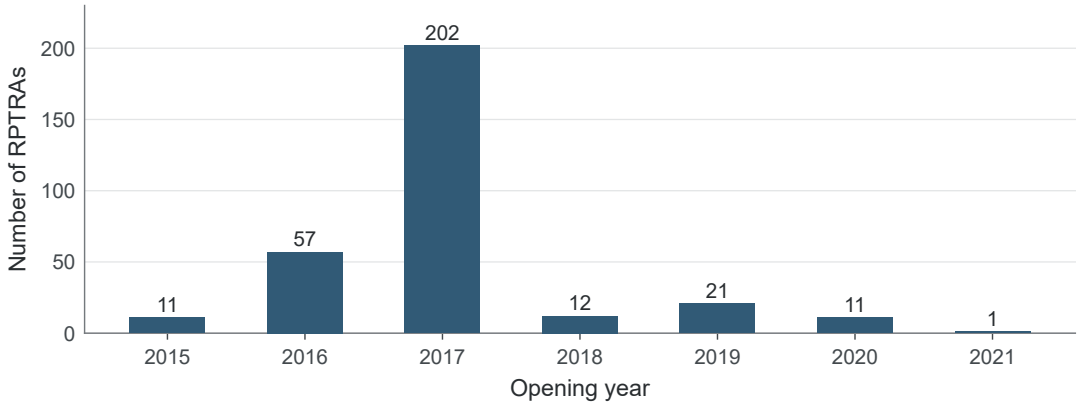
D Local Public Spaces

This appendix reports the rollout and sample diagnostics and the robustness of the street-presence estimates in Section 6. It then discusses the remaining identification concerns, documents the audit of interior photographs and reports a check on residential responses. Each exercise retains its own outcome, comparison group and inference procedure.

D.1 Programme Rollout, Coverage and Baseline Characteristics

Figure D1 reports the opening years used to define RPTRA exposure. Figure D2 shows imagery availability in the retained capture years by distance from a facility. Its denominator is the union of cells ever observed in each band, not the original sampling frame. It therefore describes coverage within the observed panel rather than the probability that an initially sampled cell enters the imagery.

Figure D1: RPTRA openings by year



Notes: Counts of the 315 RPTRA sites in the Jakarta programme inventory by administrative opening year. Bars count sites, not newly treated neighbourhoods. Jati Bersinar is dated to 14 February 2017, correcting the raw 2012 entry; 2017 therefore has 202 site openings across 91 newly treated neighbourhoods. The inventory contains openings through 2021. No sampling intervals are applicable.

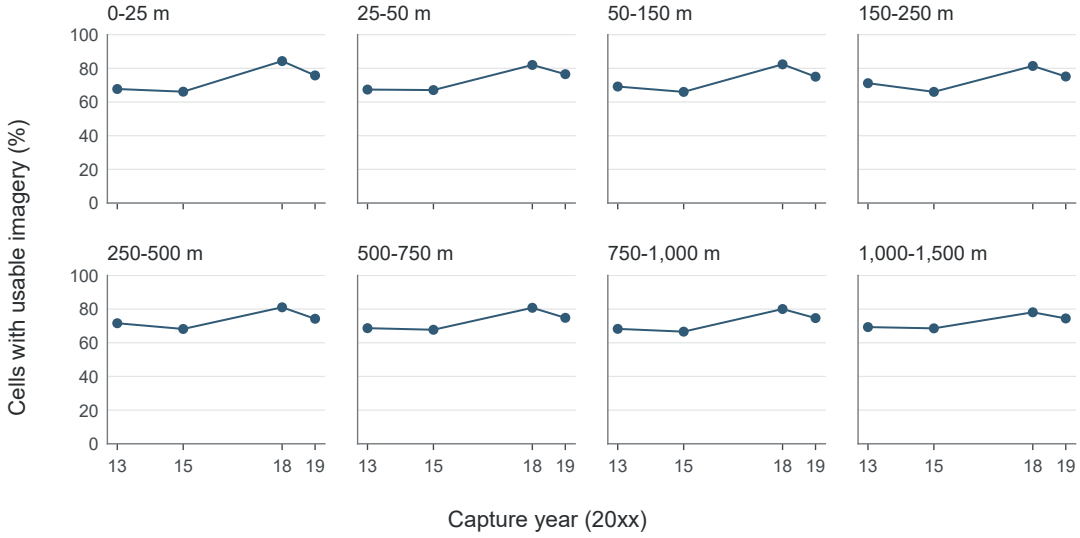
D.1.1 Baseline Balance

Table D1 describes recipient and never-treated neighbourhoods, followed by cells at different distances from RPTRA sites before opening. These are comparisons of levels. They document the samples and non-random siting, rather than testing the parallel-trends assumption.

D.2 Alternative Specifications

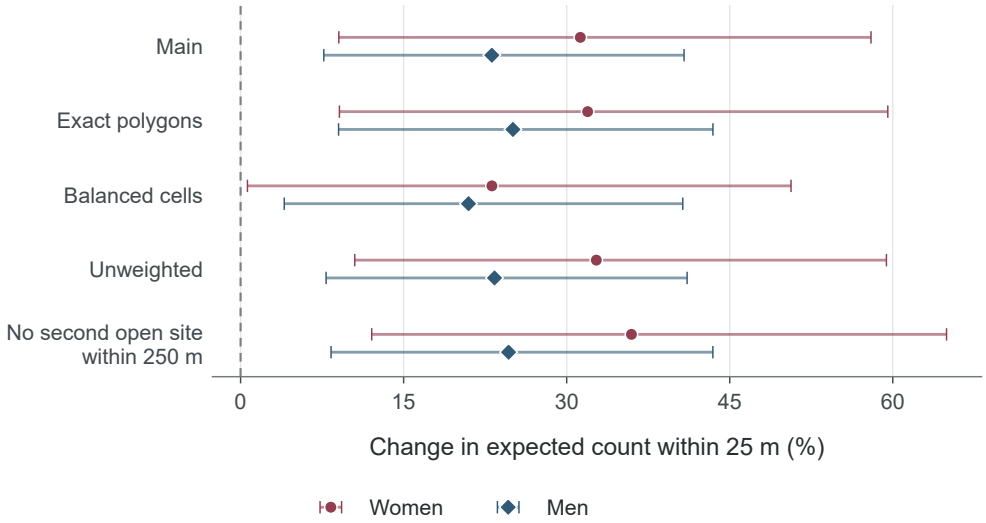
Figure D3 compares the boundary estimates under alternative geometry, sample and weighting choices. Excluding inner cells within 250 metres of another open site gives the largest female estimate among these specifications. That restriction changes the near-site sample but leaves the far comparison ring unchanged. Figure D4 uses a different outcome, the probability of observing at least one classified adult, rather than the expected count.

Figure D2: Share of cells with usable imagery, by distance band and panorama year



Notes: Jakarta Google Street View sampled 50-metre-wide hexagonal cells, assigned to the nearest RPTRA by boundary distance. In each distance band, usable cell-years are divided by the fixed union of cells ever observed in that band. Usability requires known capture month and one of 2013, 2015, 2018 or 2019. This is an unweighted descriptive coverage rate conditional on the ever-observed panel, not an estimate for the original sampling frame; no confidence intervals are shown.

Figure D3: Change in street presence within 25 metres across specifications



Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019 with known capture month. Weighted PPML uses calibrated expected adult pedestrian counts, cell, nearest-site-by-year and capture-month fixed effects, sampling and imagery-coverage weights, and nearest-RPTRA-clustered standard errors. Points report $100[\exp(\widehat{\beta}) - 1]$, with transformed pointwise 95 per cent intervals. Distances run from cell centroids to the nearest RPTRA boundary. The 0–25 metre opening interaction is compared with the 750–1,500 metre ring; all six inner-band interactions enter jointly. Rows apply the named specification or restriction: original boundary geometry, exact polygons only, balanced four-wave cells, unweighted estimation, or exclusion of inner cells within 250 metres of another open site. The unweighted row removes the weights. The last restriction applies to inner cells, not the comparison ring; other estimators are reported separately.

Table D1: Baseline characteristics of neighbourhoods and pre-opening street cells

A. Treated and never-treated neighbourhoods

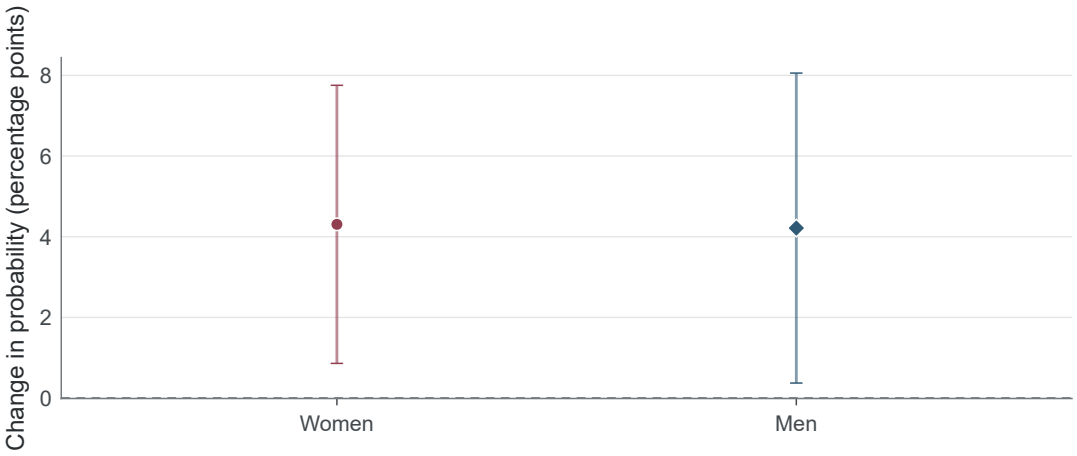
Characteristic	Treated	Never treated	Norm. diff.	N: T/C
Distance to city centre (km)	9.19	7.09	+0.456	166/95
Distance inside provincial boundary (km)	0.83	0.65	+0.361	166/95
Neighbourhood area (square km)	2.90	1.78	+0.548	166/95
Registered residents per square km, 2014	21,390	24,390	-0.183	166/95
Households per square km, 2010	4,853	5,464	-0.183	166/95
Housing floor area (square metres), 2010	71.39	72.68	-0.032	166/95
Owner-occupied housing (%), 2010	49.49	47.73	+0.173	166/95
Rented or leased housing (%), 2010	42.94	42.29	+0.059	166/95
Families, PODES 2011	10,478	6,194	+0.899	166/95
Periphery index (standard deviations)	0.02	-0.04	+0.056	166/95
Land capacity index (standard deviations)	0.14	-0.25	+0.382	166/95
Housing capacity index (standard deviations)	0.00	-0.00	+0.002	166/95
Female population share (%)	49.69	49.70	-0.006	166/95
Female labour-force participation, ages 15–64 (%)	43.17	46.75	-0.507	166/95
Male labour-force participation, ages 15–64 (%)	81.32	82.15	-0.216	166/95
Population aged 0–5 (%)	10.08	9.25	+0.523	166/95
Population aged 0–14 (%)	23.54	22.28	+0.390	166/95
Population aged 15–64 (%)	73.07	73.91	-0.329	166/95
Population aged 65+ (%)	3.39	3.81	-0.293	166/95
Members per ordinary household	3.84	3.79	+0.136	166/95
Ordinary households with a child under 15 (%)	54.22	49.82	+0.455	166/95

B. Cells observed before their assigned site opens

Characteristic	0–25 m	750–1,500 m	Norm. diff.	N: inner/far
Women per observed cell, pre-opening	0.19	0.18	+0.037	721/14,798
Men per observed cell, pre-opening	0.71	0.61	+0.090	721/14,798
Children per observed cell, pre-opening	0.20	0.13	+0.173	721/14,798
Motorcycle riders per observed cell, pre-opening	4.03	4.79	-0.155	721/14,798
Female population share (%)	50.03	49.89	+0.072	720/14,640
Female labour-force participation, ages 15–64 (%)	44.44	44.81	-0.054	720/14,640
Male labour-force participation, ages 15–64 (%)	81.70	82.45	-0.227	720/14,640
Population aged 0–5 (%)	9.88	9.96	-0.050	720/14,640
Population aged 0–14 (%)	23.02	23.30	-0.093	720/14,640
Population aged 15–64 (%)	73.61	73.51	+0.043	720/14,640
Population aged 65+ (%)	3.37	3.19	+0.136	720/14,640
Members per ordinary household	3.82	3.83	-0.044	720/14,640
Ordinary households with a child under 15 (%)	52.97	54.04	-0.125	720/14,640

Notes: Panel A gives equal weight to each neighbourhood: 166 treated by 2021 and 95 never treated. Demographic and housing measures are the supplied sampling-weighted SP2010 neighbourhood estimates; labour-force participation uses ages 15–64 in ordinary households, and household-size/child rows use ordinary households. Geometry uses the existing boundary map; the other vintages are shown in row labels. Panel B first averages each outcome within cell over usable pre-opening observations in 2013, 2015, 2018 and 2019, then uses the original cell sampling weights; all-zero cells are retained. Its demographic rows describe the cell's neighbourhood, not residents within that cell. The 159 cells with unmatched neighbourhood centroids remain in outcome rows but are excluded from demographic rows. N reports nonmissing units in each group. Normalised differences are (first-group mean minus second-group mean) divided by the square root of the mean of the two within-group variances; Panel B uses sampling-weighted means and variances. No hypothesis tests are reported.

Figure D4: Change in the probability of observing any woman or any man within 25 metres

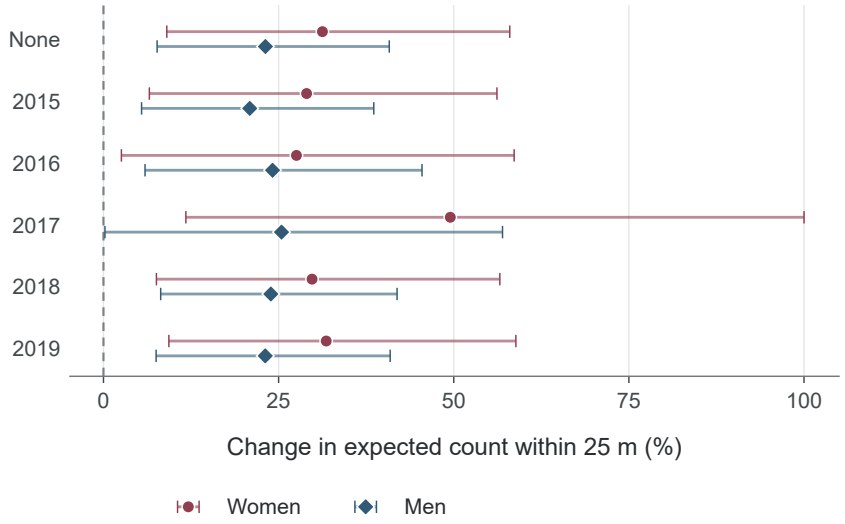


Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019. The outcome is one when at least one adult of the indicated gender is assigned under the 0.62-threshold hard classification rule. Weighted linear probability models include all six post-opening distance-band interactions, cell, nearest-site-by-year and capture-month fixed effects, and sampling and imagery-coverage weights. Points are the 0–25 metre response in percentage points relative to the omitted 750–1,500 metre boundary-distance ring. Bars are 95 per cent intervals clustered by nearest RPTRA.

D.3 Opening Cohorts and Individual Sites

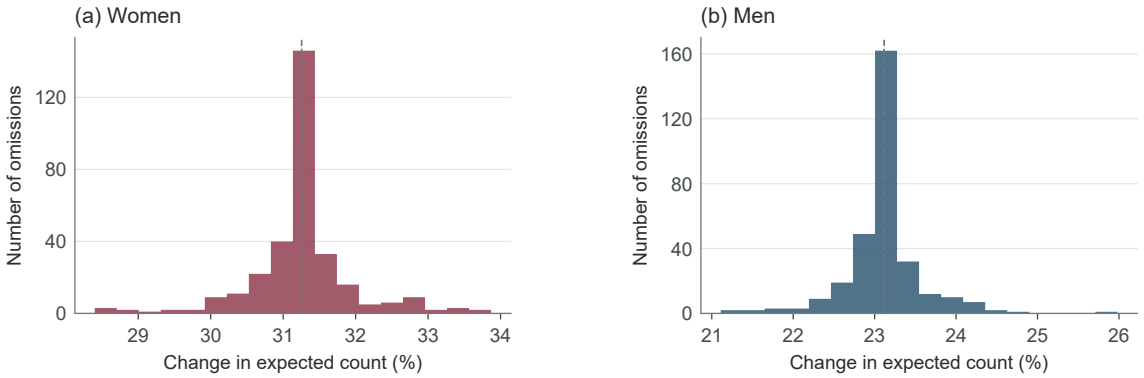
The omission exercises ask whether the main estimates depend heavily on an opening cohort or an individual site. Figure D5 removes each cohort and all cells assigned to its sites. Figure D6 does the same one site at a time. The latter histograms describe the sensitivity of the fitted coefficient to omission, not its sampling distribution.

Figure D5: Change in street presence within 25 metres, omitting each opening cohort in turn



Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019 with known capture month. Weighted PPML uses calibrated expected adult pedestrian counts, cell, nearest-site-by-year and capture-month fixed effects, sampling and imagery-coverage weights, and nearest-RPTRA-clustered standard errors. Points report $100[\exp(\widehat{\beta}) - 1]$, with transformed pointwise 95 per cent intervals. Distances run from cell centroids to the nearest RPTRA boundary. The 0–25 metre opening interaction is compared with the 750–1,500 metre ring; all six inner-band interactions enter jointly. Each omission removes all sites opened in the indicated year and every cell assigned to them, including comparison cells, before refitting. The main estimate retains all cohorts.

Figure D6: Distribution of the 0–25 metre estimate when each site is omitted in turn

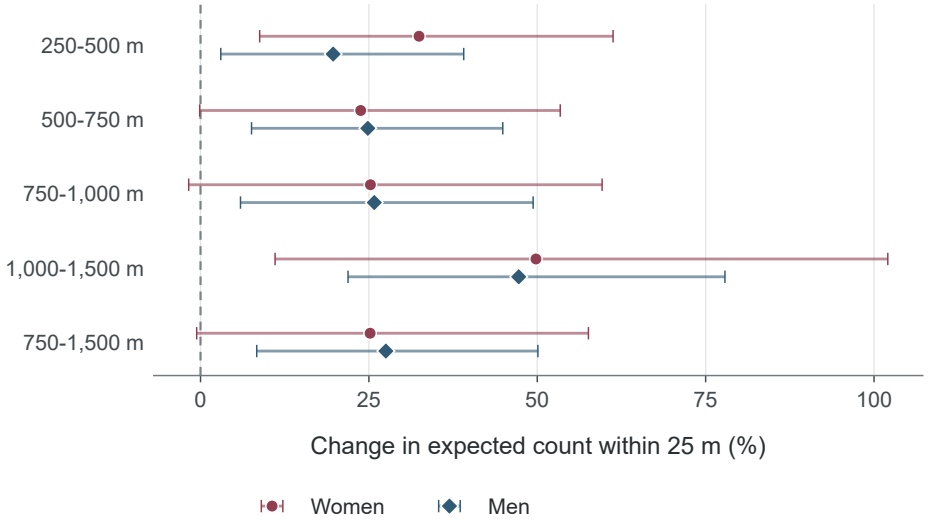


Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019 with known capture month. Weighted PPML uses calibrated expected adult pedestrian counts, cell, nearest-site-by-year and capture-month fixed effects, sampling and imagery-coverage weights, and nearest-RPTRA-clustered standard errors. The 0–25 metre opening interaction is compared with the 750–1,500 metre ring; all six inner-band interactions enter jointly. Each histogram contains 314 refits, each omitting one retained site and all cells assigned to it. Women are on the left and men on the right. The horizontal axis is $100[\exp(\widehat{\beta}) - 1]$, and the vertical line is the corresponding full-sample estimate. Histogram frequencies describe sensitivity to site omission, not a sampling distribution or confidence interval.

D.4 Alternative Spatial Comparisons

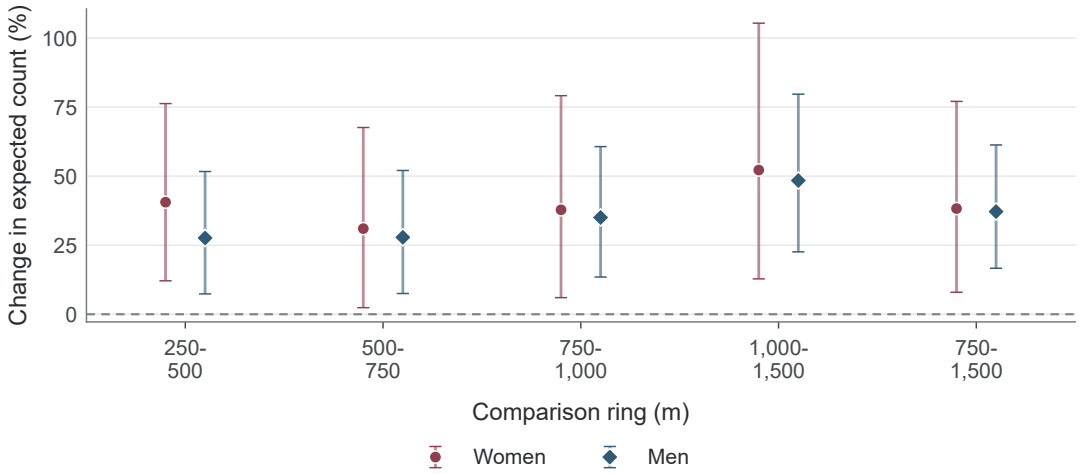
The following checks change either the comparison ring or the radius of the treated area. Figure D7 holds the treated band at 0–25 metres and changes the single comparison ring, allowing the retained site set to vary. Figure D8 repeats the comparison on the common set of 111 sites with support in every component ring. This separates changes in the spatial comparison from changes in the site set.

Figure D7: Change in street presence within 25 metres, by choice of single comparison ring



Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019 with known capture month. Weighted PPML uses calibrated expected adult pedestrian counts, cell, nearest-site-by-year and capture-month fixed effects, sampling and imagery-coverage weights, and nearest-RPTRA-clustered standard errors. Points report $100[\exp(\widehat{\beta}) - 1]$, with transformed pointwise 95 per cent intervals. Distances run from cell centroids to the nearest RPTRA boundary. Each regression compares cells within 25 metres with the single, non-nested ring named on the axis, excluding intervening cells. Retained site sets may differ across regressions; the common-site comparison is reported separately.

Figure D8: Change in street presence within 25 metres by comparison ring, common set of sites



Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019 with known capture month. Weighted PPML uses calibrated expected adult pedestrian counts, cell, nearest-site-by-year and capture-month fixed effects, sampling and imagery-coverage weights, and nearest-RPTRA-clustered standard errors. Points report $100[\exp(\widehat{\beta}) - 1]$, with transformed pointwise 95 per cent intervals. Distances run from cell centroids to the nearest RPTRA boundary. Each regression compares cells within 25 metres with the named comparison ring. All five regressions retain the same 111 sites with positive paired support and actual PPML support in every component ring. Two initially eligible sites lack estimation support in at least one ring. This restriction holds the site set fixed across spatial comparisons.

Two additional specifications retain the fixed effects, weights and clustering of the main model. The conditional expectations below condition on the cell, assigned-site-by-year and capture-month effects, together with the stated opening interactions. The first replaces the six band interactions with one cutoff interaction:

$$E[y_{itm} \mid \cdot] = \exp\Big(\alpha_i + \delta_{p(i)t} + \mu_m + \beta^{(c)}\,\mathbf{1}[d_i \le c]\,D_{itm}\Big), \tag{D1}$$

I estimate (D1) separately for $c = 25$, 50 and 250 metres. Each regression retains cells within c metres and comparison cells 750–1,500 metres from their nearest site. Cells between c and 750 metres are excluded, so the comparison ring is the same for every cutoff.

The second specification follows the 2017 opening cohort, which contains 202 of the 315 sites:

$$E[y_{itm} \mid \cdot] = \exp\Big(\alpha_i + \delta_{p(i)t} + \mu_m + \textstyle\sum_{k\in\{-4,1,2\}} \gamma_k\,\mathbf{1}[d_i \le 25]\,\mathbf{1}[t - 2017 = k]\Big), \tag{D2}$$

The sample contains cells within 25 metres of 2017 sites and the 750–1,500 metre comparison ring around those sites. Event time is capture year minus 2017. The retained years correspond to $k = -4, -2, +1, +2$, with $k = -2$ (2015) omitted. Each γ_k measures the change in log expected occupancy in the inner band from 2015 to $2017 + k$, relative to the change in the comparison ring.

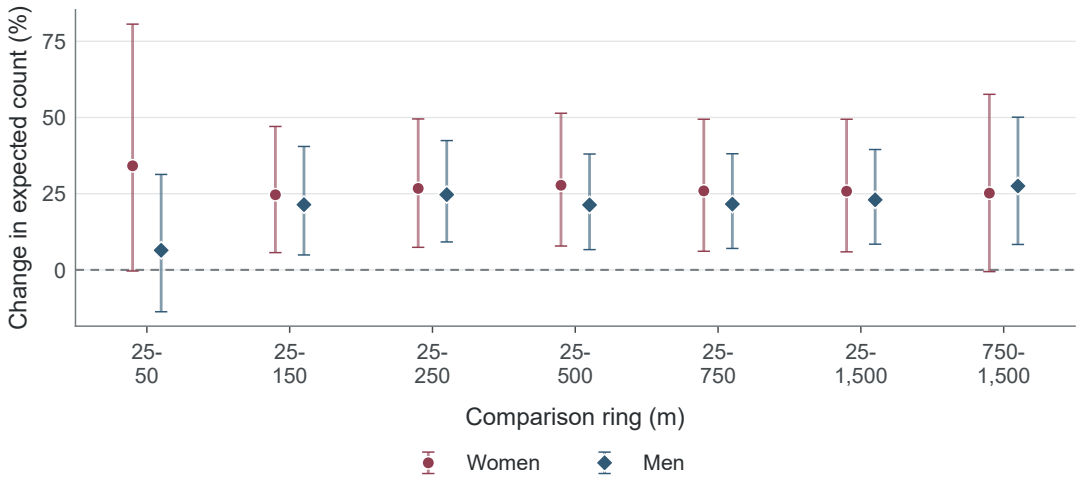
The four capture years provide limited information about dynamics before and after opening. The within-site spatial comparison and site-by-year effects account for permanent street differences and shocks shared within each catchment, but do not replace evidence on differential pre-trends. Figure D2 reports the remaining coverage differences; it does not establish equal annual selection across bands. Table D2 gives the cutoff estimates and Figure D9 widens the comparison ring while holding the near band fixed.

Table D2: Change in street presence at alternative distance cutoffs

	Within 25 m	Within 50 m	Within 250 m
Women	25.2* [-0.6, 57.6]	14.8 [-5.9, 39.9]	2.8 [-6.5, 13.0]
Men	27.5*** [8.3, 50.1]	24.7*** [8.9, 42.8]	-0.6 [-7.2, 6.4]
Children	29.3 [-5.5, 76.8]	29.1** [0.2, 66.2]	-2.7 [-13.3, 9.2]
All pedestrians	27.5*** [7.7, 50.9]	24.0*** [8.3, 42.0]	-0.1 [-6.7, 7.0]
Motorcycle riders	9.5* [-1.3, 21.5]	5.9 [-1.5, 13.9]	0.0 [-3.6, 3.8]

Notes: Jakarta Google Street View, 2013, 2015, 2018 and 2019; machine-expected pedestrian, gender, child and rider counts use sampling and imagery-coverage weights. Each column estimates Equation (D1): treated cells lie within the indicated distance of their nearest site, the comparison group is the 750–1,500 metre ring around the same sites, and cells in between are excluded. Entries are percentage changes from weighted Poisson estimates with 95 per cent confidence intervals in brackets. Cell, park-by-year and capture-month fixed effects are included throughout. Standard errors are clustered at the nearest park. Samples differ across rows because of Poisson separation: within 25 metres the pedestrian, female, male and child rows use 33,080 observations across 269 sites and the rider row 46,760 across 287. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure D9: Change in street presence within 25 metres, by width of the comparison ring

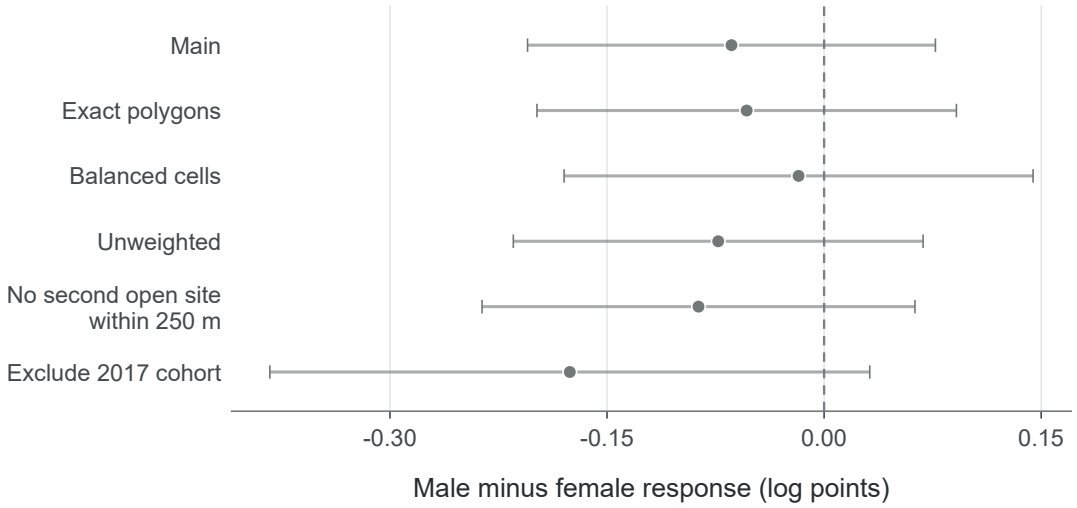


Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019 with known capture month. Weighted PPML uses calibrated expected adult pedestrian counts, cell, nearest-site-by-year and capture-month fixed effects, sampling and imagery-coverage weights, and nearest-RPTRA-clustered standard errors. Points report $100[\exp(\widehat{\beta}) - 1]$, with transformed pointwise 95 per cent intervals. Distances run from cell centroids to the nearest RPTRA boundary. Each regression holds the treated band at 0–25 metres and uses the comparison interval named on the axis, excluding all other cells. The comparison ring widens across specifications; single non-nested rings are shown separately.

D.5 Gender Differences, Alternative Estimators and Classification

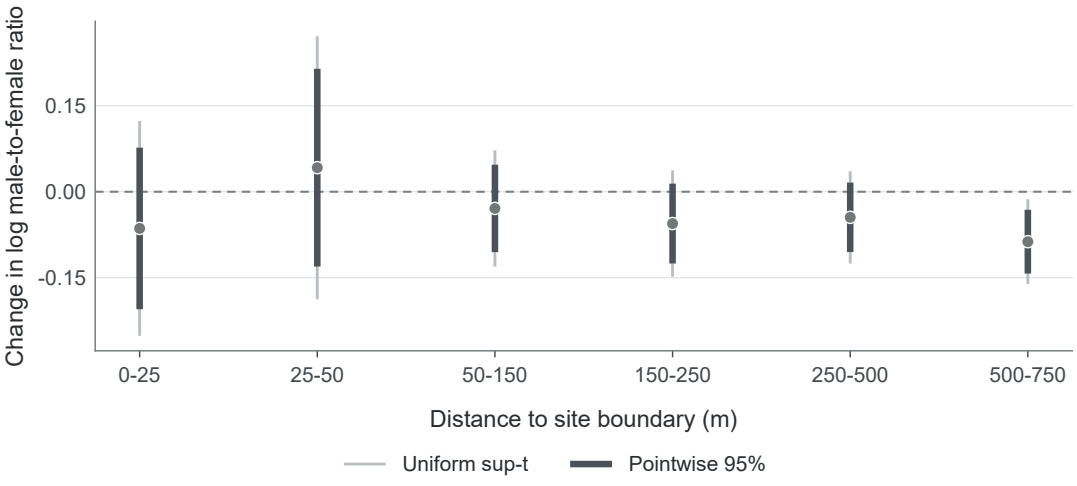
Figure D10 compares the male-minus-female boundary coefficient across specifications. Figure D11 reports that contrast for each distance band, retaining the covariance between gender coefficients. Table D3 then uses the alternative staggered Poisson estimator. Its level effects and pre-opening-mean normalisation differ from the log responses in the main profile. Finally, Figures D12 and D13 examine alternative classification rules and the resulting probability distributions.

Figure D10: Male-minus-female difference in the response within 25 metres, across specifications



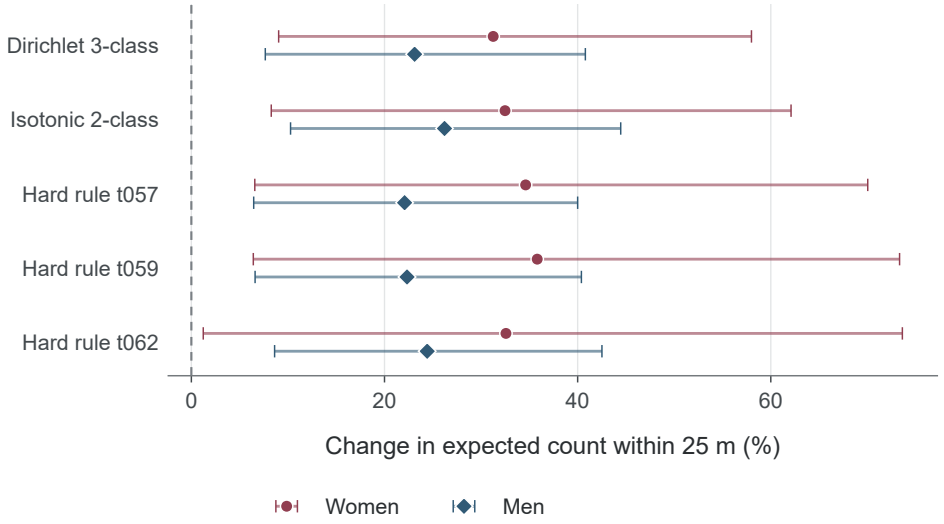
Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019. Stacked PPML uses calibrated male and female adult pedestrian counts with gender-specific cell, nearest-site-by-year and capture-month fixed effects, plus sampling and coverage weights except in the unweighted row. Points are male-minus-female differences in the 0–25 metre log response, not percentage-point differences between transformed effects. All six post-opening distance-band interactions are included relative to the 750–1,500 metre ring. Bars are pointwise 95 per cent intervals clustered by site. Each row uses its named restriction; the last omits the 2017 cohort.

Figure D11: Difference between the male and female response, by distance from the site boundary



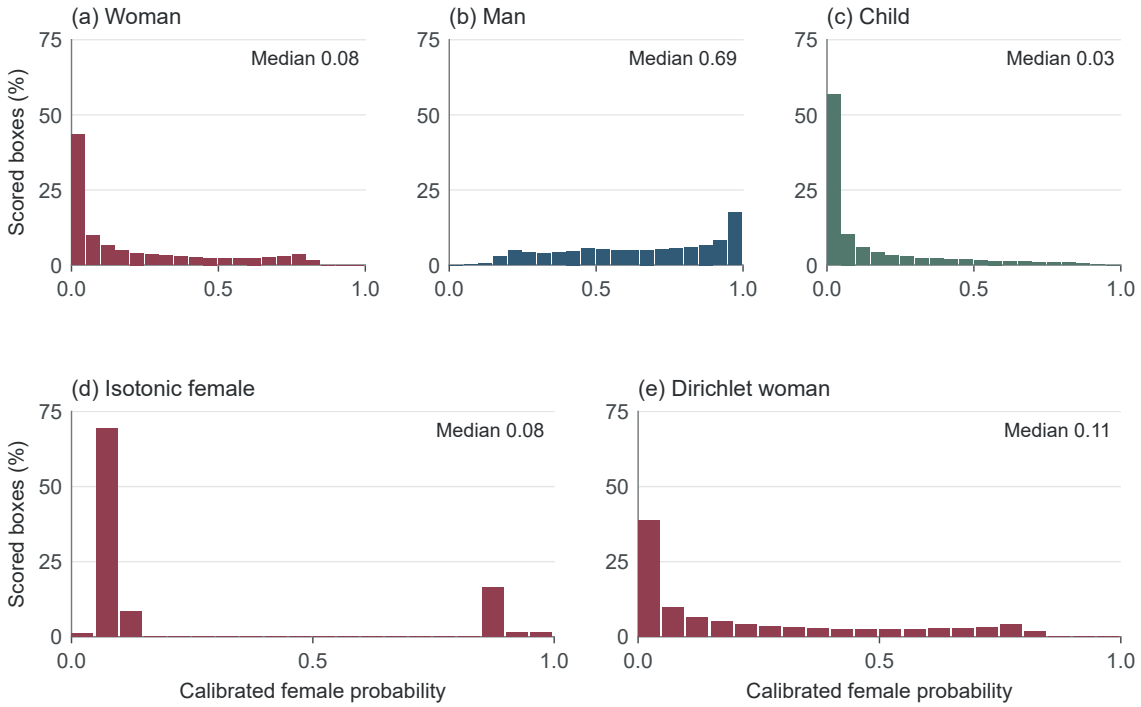
Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019. Stacked PPML models calibrated male and female adult pedestrian counts with gender-specific cell, nearest-site-by-year and capture-month fixed effects and sampling and coverage weights. Points are male-minus-female differences in post-opening log responses for the six boundary-distance bands, relative to the omitted 750–1,500 metre ring. Joint estimation retains the covariance between gender coefficients. Inner bars are pointwise 95 per cent intervals; outer bars are uniform sup-t bands across the six contrasts, with critical value 2.61. Standard errors cluster by nearest RPTRA.

Figure D12: Change in street presence within 25 metres under alternative classification arms



Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019. Separate weighted PPML fits use expected adult pedestrian counts from three-class Dirichlet calibration or isotonic interpolation, or counts under the indicated hard-label decision rules. All include six post-opening distance interactions, cell, nearest-site-by-year and capture-month fixed effects, sampling and coverage weights, and site-clustered standard errors. Points give $100[\exp(\widehat{\beta}) - 1]$ within 25 metres relative to the 750–1,500 metre ring, with transformed 95 per cent intervals. Poisson separation leaves 54,029–61,525 female observations under hard rules versus 118,588 under Dirichlet calibration; these are not identical estimation samples.

Figure D13: Distribution of calibrated class probabilities and comparison of calibration methods



Notes: Google Street View pedestrian detections from the Jakarta classifier audit. Panels (a)–(c) show the distributions of Dirichlet three-class probabilities across 235,818 scored detections. Panels (d)–(e) compare female probability distributions under the Dirichlet and isotonic calibrations on the same 203,602 detections scored by both methods. Distributions are unweighted, with bins of width 0.05; unscored boxes are excluded.

Table D3: Staggered Poisson level effects and their normalisation

	Level effect (SE)	p	Pre-opening mean, treated	Implied per cent [95% CI]
Women	0.058 (0.022)	0.008	0.202	28.9 [7.6, 50.2]
Men	0.154 (0.050)	0.002	0.725	21.3 [7.8, 34.7]

Notes: Jakarta Google Street View, 2013, 2015, 2018 and 2019; calibrated expected adult pedestrian counts in the 0–25 metre boundary band and 750–1,500 metre comparison ring. The staggered Poisson estimator reports average treatment effects on treated cells in counts per cell, using sampling/coverage weights and nearest-site clustering, on 54,312 observations. Parentheses are standard errors. Treatment uses a four-wave index: 32 observations in the opening wave but captured before opening are excluded; later observations remain. The estimator includes cell and time structure but omits nearest-site-by-year effects. Implied per cent divides the level effect and its 95 per cent interval by the fixed, weighted pre-opening mean for treated cells in this sample. It is not the main distance-profile log response; inclusion of all-zero cells also changes the normalising mean.

D.6 Identification, Remaining Limitations and Inference

For the street design, identification requires the same proportional counterfactual change in the near and comparison bands. Fixed effects remove permanent spatial differences and common catchment-year shocks, but not every opening-related change. The main concerns are site selection, coincident timing, spillovers and measurement.

Non-random park placement. The city selected RPTRA sites. Permanent differences between selected and comparison locations are not themselves violations of the street design; selection matters if it predicts different changes around opening. Recipient neighbourhoods have larger populations and areas and differ in some amenity densities. Female population shares are 49.69 and 49.70 per cent in recipient and never-treated neighbourhoods, and ordinary-household sizes are 3.84 and 3.79. Female labour-force participation is 43.17 versus 46.75 per cent; households with a child under fifteen account for 54.22 versus 49.82 per cent. Table D1 gives the full comparison and pre-opening cell statistics.

Officials described land availability as central to siting, including parcels already in public ownership or informal public use. That account does not make the sites exogenous. Within neighbourhoods, lower-valued parcels may follow different trajectories. Where construction involved clearance, the observed change can combine population displacement and amenity provision; the analysis does not separate them.

Staggered timing. With staggered adoption and heterogeneous effects, a two-way fixed-effects coefficient need not be a convex average of treatment effects because already-treated units can contribute as comparisons (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Borusyak et al., 2024). In the street design, site-by-year effects make the main contrast a change in the spatial gradient within each catchment. That structure does not remove every concern about heterogeneous-effect weighting or common calendar-year changes in near–far gradients.

The alternative Poisson estimator of Wooldridge (2023) groups sites by their first post-opening panorama wave and allows group-by-wave effects, using untreated cells as comparisons without site-by-year effects. It gives average effects on treated cells of 0.058 women and 0.154 men. These equal 28.9 and 21.3 per cent of the corresponding treated pre-opening means (Table D3); they are normalised level effects, not the main profile coefficients.

Leaving out each opening cohort gives female estimates of 27.5–49.5 per cent and male estimates of 20.9–25.4 per cent. The sign does not depend on one cohort, but the female magnitude varies more.

Omitting each of the 314 sites with all its assigned cells gives ranges of 28.4–33.9 per cent for women and 21.1–26.0 per cent for men (Figure D6).

Spillovers and the ring design. An effect reaching the comparison ring is subtracted from the near-site response. A positive comparison-ring response can therefore reduce the estimated positive contrast (Butts, 2023; Clarke, 2017). Comparison cells are assigned at least 750 metres from a site, and the reported profile does not reject zero jointly beyond 250 metres. Failure to reject does not establish an absence of spillovers.

Interpreting each band as a causal spatial effect also requires its untreated trend to match the far comparison group's trend. Ring sensitivity can reflect unequal counterfactual trends, treatment effects reaching the comparison ring, or both. It does not identify which explanation applies. Excluding inner cells within 250 metres of a second open site increases the female estimate to 36.0 per cent and the male estimate to 24.7 per cent. Because this restriction changes only the inner sample, it is not a test of contamination in the far ring.

Other-site exposure changes during the same rollout. Among cell-years within 25 metres of the assigned site, 6.3 per cent have a second open site within 250 metres and 53.3 per cent have one within 1.5 kilometres. The corresponding shares within 1.5 kilometres are 45.0 and 27.9 per cent in the two comparison bands. Among 615 near-site cells observed before and after opening, the mean number of other open sites within 1.5 kilometres rises from 0.18 to 4.48. This calculation averages within each cell and period, then gives equal weight to cells. In the retained imagery years, 2016 and 2017 openings are both first observed in 2018, limiting timing variation that could separate their effects.

Adding time-varying counts of other open sites within 250 metres and 1.5 kilometres gives near-site estimates of 30.7–32.2 per cent for women and 23.0–23.7 per cent for men. The 1.5-kilometre count coefficient is close to zero for both. The 250-metre count coefficient is -0.118 for women ($p = 0.09$) and -0.073 for men, consistent with the larger estimate after excluding doubly exposed cells.

A further specification adds eight opening-by-distance indicators for the closest non-assigned site, retaining the primary assignment and fixed effects. It gives 33.5 per cent for women and 25.9 per cent for men on 118,584 observations and 314 site clusters (Table D5).[20] This allows distance-specific exposure to one other site, not unrestricted effects of every nearby facility. The design remains a local comparison within assigned catchments and does not identify net movement across the city.

Functional form and inference. Alternative outcomes also give positive boundary estimates. The inverse-hyperbolic-sine specification gives coefficients of 0.049 for women and 0.077 for men. Linear probability models using the 0.62 hard-label rule give 4.3 and 4.2 percentage points, with standard errors of 1.8 and 2.0 (Figure D4). These outcomes have different scales and do not estimate an identical parameter.

Across the main model, exact polygons, balanced four-wave cells, unweighted estimation and the overlap restriction, female estimates range from 23.1 to 36.0 per cent and male estimates from 21.0 to 25.1 per cent. Each is positive and significant at five per cent; the main female estimate is 31.3 per cent.

[20] Distances are recomputed from projected cell-polygon centroids to the same site polygons, with the original circular fallback. In 19 of 63,033 cells, the fixed assignment differs from the geometric minimum by at most 1.40 metres. The check selects the closest site other than the assigned one and leaves the primary assignment unchanged.

Clustering at the sub-district level gives standard errors of 0.097 for women and 0.066 for men across 42 clusters. This is a coarser-clustering check, not a spatial HAC estimator.

Because those clusters are few and unbalanced, I also use a null-imposed Poisson cluster-score bootstrap with Rademacher weights and 9,999 draws. The p-values are 0.008 for women, 0.005 for men and 0.374 for the male-minus-female contrast (Table D4).

Measurement error in the outcome. Detection and classification are separate sources of error. The interior audit misses 46.4 per cent of hand-labelled women and 34.1 per cent of men, but those selected photographs do not establish street-panel error rates. The relevant question is whether eligible people are detected and classified comparably in near and far streets before and after opening. That has not been tested directly. An opening-related change in crowding or visibility could alter counts even without a change in the underlying occupancy of the same magnitude. Positive results across calibrated, isotonic and hard-label outcomes test class allocation on the detected sample, not detection stability (Figure D12).

Table D4: Distance-profile tests and resampling checks

Test or quantity	Women	Men	Male-minus-female
0-25 m = 25-50 m: p	0.0678	0.2424	-
250-750 m jointly zero: p	0.1210	0.1802	-
50-750 m jointly zero: p	0.3551	0.3985	-
Gender differences, bands 3-6 jointly zero: p	-	-	0.0416
Boundary minus mean bands 3-6 gender difference: p	-	-	0.8882
Null-imposed cluster-score bootstrap: p	0.0084	0.0053	0.3743
Whole-park omission range (%)	28.3989 to 33.8693	21.1103 to 25.9811	-

Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019. The outcome is calibrated expected adult pedestrian counts. Weighted PPML includes cell, nearest-site-by-year and capture-month fixed effects and sampling and imagery-coverage weights. All six post-opening distance-band interactions enter relative to the omitted 750–1,500 metre boundary-distance ring. Band and gender-difference Wald tests use site-clustered covariance; stacked gender tests use gender-specific fixed effects. The outer-band mean equally weights 50–150, 150–250, 250–500 and 500–750 metre coefficients. The null-imposed, nuisance-adjusted PPML cluster-score bootstrap uses 42 sub-district clusters and 9,999 Rademacher draws, with a plus-one p-value correction; it does not refit the nonlinear model. Site-omission ranges use 314 complete refits per gender, each removing one site and its comparison cells.

Table D5: Street presence within 25 metres with other-site exposure controls

Specification	Women: % [95% CI]	Men: % [95% CI]	Rows	Parks
Main profile	31.3 [9.0, 58.0]	23.1 [7.7, 40.8]	118588	314
Other open sites: 250 m	32.2 [9.7, 59.3]	23.7 [8.2, 41.4]	118588	314
Other open sites: 1500 m	30.7 [8.7, 57.2]	23.0 [7.4, 40.9]	118588	314
Both other-site counts	31.6 [9.3, 58.5]	23.6 [7.9, 41.5]	118588	314
Exclude ever overlap 250 m	36.0 [12.1, 64.9]	24.7 [8.3, 43.5]	115257	311
Second-site distance-by-opening bands	33.5 [10.6, 61.2]	25.9 [9.8, 44.4]	118584	314

Notes: Jakarta Google Street View cell-year panel, 2013, 2015, 2018 and 2019. The outcome is calibrated expected adult pedestrian counts. Weighted PPML includes cell, nearest-site-by-year and capture-month fixed effects and sampling and imagery-coverage weights. Entries report $100[\exp(\widehat{\beta}) - 1]$ for 0–25 metres relative to 750–1,500 metres, with transformed pointwise 95 per cent intervals and site-clustered standard errors. All six nearest-site distance interactions remain. Count controls exclude the assigned site. The exclusion removes all usable observations of cells ever exposed to another open site within 250 metres. The final row adds eight opening-by-distance indicators for the nearest non-assigned site, with bands ending at 1,500 metres; distances use projected cell-polygon centroids and the same site polygons/circular fallback. Original nearest-site assignment remains fixed.

D.7 Interior Photographs

The interior audit labels 343 photographs from 208 sites, using the stratified design to weight the audit to the full photo collection. It compares machine output with human labels rather than treating machine composition as ground truth. Because the classifier was trained on street imagery, its performance need not transfer to ordinary interior photographs.

The pipeline counts 52.7 per cent as many people as the weighted human estimate. Among estimated real people, 44.6 per cent were never detected, 13.5 per cent were detected but removed by the rider gate, 0.5 per cent were removed by the vehicle-overlap rule, and 41.3 per cent were counted. In the audited sample, 21.5 per cent of pre-gate detections matched no hand-labelled person. False positives can therefore offset missed people in aggregate totals; similar adult totals are not evidence of correct person-level matching.

Age classification is also problematic in these photographs. The hard-label file assigns no children, while children constitute 44.7 per cent of the weighted hand count. Calibrated child probability mass recovers only 7.2 per cent of the hand child estimate. Non-detection differs by gender: 46.4 per cent of hand-labelled women and 34.1 per cent of men were never detected. Grouping may contribute, but these interior diagnostics do not establish the error rates or their causes in the street panel.

Table D6 compares blind human labels with machine counts and composition across the interior-photo categories. The intervals reflect audit sampling, not upload selection or annotation error.

Table D6: Hand audit of interior photographs, by use category

Photo category	Photos	Machine/hand	Hand female	Calibrated female	Hand child
Ordinary: people	142	66.7 [60.1, 73.8]	47.4 [37.0, 57.9]	48.7	63.3 [54.5, 71.2]
Ordinary: no pipeline people	75	49.0 [32.7, 78.6]	37.4 [4.7, 56.9]	53.0	22.9 [10.0, 40.7]
Organised event	68	42.6 [35.2, 51.5]	73.6 [65.1, 80.9]	60.2	35.5 [25.9, 45.0]
Posed/group photo	58	59.0 [43.3, 76.7]	80.4 [68.1, 87.9]	53.5	27.4 [12.4, 44.5]

Notes: Google Maps interior photographs from 208 Jakarta RPTRAs. The 343 blind hand-labelled photographs were sampled in six strata and weighted to the 2,436-photograph run. Machine-to-hand is 100 times weighted hard-detection counts divided by weighted hand counts. Female and child shares are percentages; the child denominator is all people, while female shares use identified adults. Brackets on the count ratio and the hand-labelled female and child shares are 95 per cent stratified image-bootstrap intervals from 1,999 draws (seed 20260922); the all-gated census stratum is held fixed. Female adult shares exclude children and unknown hand labels. The calibrated column sums the three-class probabilities and is not a hard-label share. Intervals describe sampling variation under the audit design, not uncertainty in hand annotation.

Table D7 retains the earlier hand-labelled interior versus hand-labelled street diagnostic, separating ordinary-use, all-photo, event and posed-photo comparisons. Its nearby columns use distance to park location pins. This differs from the main comparison in Table 4, which uses machine-classified exterior street cells and distance to park boundaries. Nearby columns match interiors with post-opening streets around the same parks, while distant and citywide columns use broader benchmarks. Different columns can retain different parks and photographs. These are descriptive composition comparisons, not effects on representative attendance.

D.8 Registered Population

If residents value RPTRAs, households may move towards them, and a residential response could contribute to the street estimates. I examine registered population in the 261 mainland neighbourhoods in 2014, 2015, 2016, 2018, 2020 and 2021, using the doubly robust estimator of Sant'Anna and Zhao (2020)

Table D7: Interior female composition across photograph samples and street comparison groups

	Within 100 m Post-opening	Within 250 m Post-opening	Beyond 250 m All dates	Beyond 250 m Post-opening	Citywide All dates
Same-park restriction	Yes	Yes	No	No	No
A. Ordinary-use interior photographs					
Interior women (%)	54.2	48.8	46.7	46.7	46.7
Street women (%)	30.3	24.6	24.3	25.1	24.7
Difference (pp)	24.0	24.2	22.4	21.7	22.1
95% interval	[-24.0, 71.9]	[12.0, 36.4]	[12.8, 32.1]	[12.1, 31.3]	[12.5, 31.7]
p-value	0.298	< 0.001	< 0.001	< 0.001	< 0.001
Images: interior / street	17 / 34	123 / 295	217 / 9,406	217 / 6,039	217 / 10,216
Parks: interior / street	15 / 15	88 / 88	161 / 253	161 / 250	161 / 263
Inference clusters	13	77	264	262	272
B. All hand-labelled interior photographs					
Interior women (%)	66.9	63.1	68.0	68.0	68.0
Street women (%)	26.1	25.1	24.3	25.1	24.7
Difference (pp)	40.9	38.1	43.7	43.0	43.4
95% interval	[15.1, 66.7]	[27.3, 48.9]	[37.1, 50.3]	[36.3, 49.6]	[36.8, 49.9]
p-value	0.004	< 0.001	< 0.001	< 0.001	< 0.001
Images: interior / street	24 / 47	184 / 368	343 / 9,406	343 / 6,039	343 / 10,216
Parks: interior / street	20 / 20	112 / 112	208 / 253	208 / 250	208 / 263
Inference clusters	18	102	276	275	281
C. Organised-event interior photographs					
Interior women (%)	80.0	68.9	73.6	73.6	73.6
Street women (%)	16.7	26.2	24.3	25.1	24.7
Difference (pp)	63.3	42.7	49.2	48.5	48.9
95% interval	[-83.1, 209.7]	[25.9, 59.5]	[41.1, 57.4]	[40.3, 56.7]	[40.8, 57.1]
p-value	0.115	< 0.001	< 0.001	< 0.001	< 0.001
Images: interior / street	2 / 8	31 / 81	68 / 9,406	68 / 6,039	68 / 10,216
Parks: interior / street	2 / 2	28 / 28	62 / 253	62 / 250	62 / 263
Inference clusters	2	28	260	257	266
D. Posed/group interior photographs					
Interior women (%)	62.5	68.7	80.4	80.4	80.4
Street women (%)	18.9	23.8	24.3	25.1	24.7
Difference (pp)	43.6	44.9	56.0	55.3	55.7
95% interval	[2.0, 85.2]	[25.4, 64.4]	[46.2, 65.9]	[45.4, 65.2]	[45.8, 65.6]
p-value	0.045	< 0.001	< 0.001	< 0.001	< 0.001
Images: interior / street	5 / 11	30 / 79	58 / 9,406	58 / 6,039	58 / 10,216
Parks: interior / street	4 / 4	23 / 23	43 / 253	43 / 250	43 / 263
Inference clusters	4	23	256	254	265

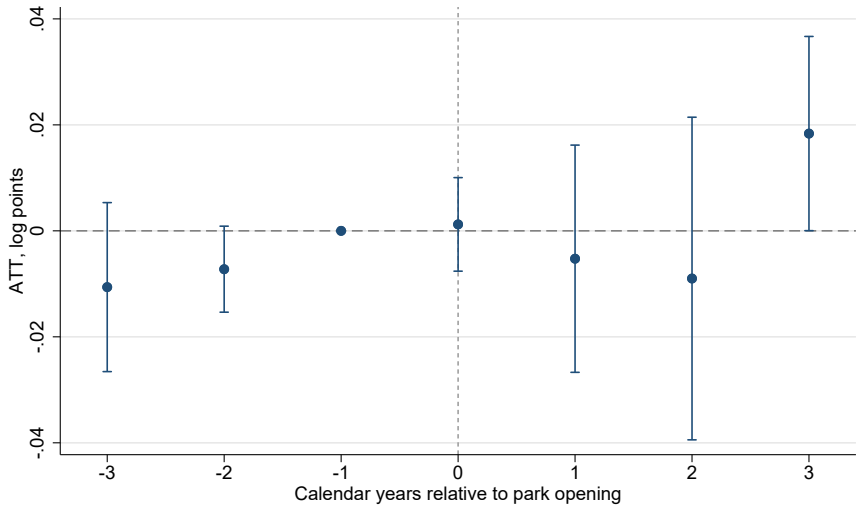
Notes: Women as a share of hand-identified adults; children and unidentified adults are excluded from that denominator. Google Maps interior photographs retain supplied inverse-probability weights; Jakarta Street View hand counts are unweighted and include separately coded non-motorised riders. Nearby columns restrict both samples to the same parks, separately for each interior category and distance window. Distant and citywide columns are broader benchmarks without a common-park restriction. All-date street coverage is 2013–2024 excluding 2020–2021; interior dates are not aligned. Distances are to park pins, not boundaries. Panel A includes both ordinary-use strata, including photographs without pipeline detections; panel B includes all 343 hand-labelled photographs before any common-park restriction. Differences subtract street from interior shares in percentage points. The 95 per cent intervals and two-sided tests cluster by park, allowing covariance where parks overlap. Reported p-values are unadjusted; Holm-adjusted values for the full 25-test diagnostic family (including the ordinary-with-detections sensitivity) are supplied in the data. Sample park counts include parks without identified adults. Event and posed 100 m comparisons have only two and four inference clusters; their inferential results are especially unreliable. The columns overlap and are not independent. Selection and annotation uncertainty are not corrected; these are descriptive comparisons, not causal attendance effects.

for staggered first openings (Callaway and Sant'Anna, 2021). The outcome is log registered population, and comparisons use never-treated and not-yet-treated neighbourhoods with distance to the provincial boundary as covariates.

The positive estimates arise in 2020 and later, after the street-imagery panel ends. Pooling the eligible cohorts, the estimates are -0.53 and -0.90 per cent one and two years after opening, and 1.84 per cent three years after opening (Figure D14). For the 2017 cohort alone, the estimates relative to 2016 are -1.14 per cent in 2018, 2.01 per cent in 2020 and 2.80 per cent in 2021. The log female-to-male population ratio does not change detectably.

The later estimates coincide with a citywide break in the register. Registered population grows by 5.34 per cent a year between 2018 and 2020, and by 2021 the register exceeds the 2020 census long-form benchmark by 5.5 per cent. Differential registration changes could therefore account for part of the later estimates. The timing nonetheless differs from that of the street response, which the 2013–2019 panel records before these population changes. The model in Section 7 holds residence fixed.

Figure D14: Neighbourhood population by years relative to first park opening, eligible cohorts



Notes: Jakarta civil-register population in 2014, 2015, 2016, 2018, 2020 and 2021. Treatment is the first RPTRA opening in a neighbourhood. Doubly robust Callaway–Sant'Anna estimates pool the 2015, 2016, 2017 and 2019 cohorts, using never-treated and not-yet-treated eligible neighbourhoods as comparisons. Covariates are standardised distance to the provincial boundary and its square. The outcome is the log of total registered population; points are log changes, so multiplying by 100 gives approximate percentage changes. Event time is calendar year minus first-opening year, with -1 omitted. Cohort weights reflect treated-neighbourhood counts at each event time. Bars are Bonferroni simultaneous 95 per cent intervals across the six reported event times within each series, using influence-function standard errors that allow within-neighbourhood dependence. These are administrative counts, not enumerated residents.

D.9 Assessed Land Values

The 2015 and 2019 ZNT layers supply repeated cross-sections of administrative assessment zones in North and South Jakarta. They are not transaction prices, zone boundaries change between vintages, and the supplied South Jakarta layer retains duplicated polygons. Distances are measured to the recorded pins of sites opened in 2016–2018, using the corrected Jati Bersinar date. Thus these checks have different geometry and support from the street-panel design.

The rerun 250-metre within-neighbourhood contrast is -0.047 log points (standard error 0.034,

$p = 0.176$). Restricting to North Jakarta gives 0.038, while weighting by zone area gives 0.046. The smaller-radius contrasts are more negative. These unstable signs and the two-vintage administrative measurement do not establish a causal capitalisation effect or rule out residential sorting.

Table D8: Assessed land-value contrasts near RPTRAs, 2015 and 2019

Sample / radius	Log contrast	SE	p	Zones
North and South, 250 m	-0.047	0.034	0.176	3,797
North only, 250 m	0.038	0.068	0.584	691
Area weighted, 250 m	0.046	0.091	0.615	3,797
Certain pins, 250 m	-0.040	0.036	0.262	3,683
Exclude 2015-site controls	-0.043	0.035	0.223	3,790
Controls beyond 500 m	-0.015	0.051	0.770	4,750
Future-site controls	0.021	0.262	0.937	588
Within-neighbourhood, 100 m	-0.119	0.050	0.024	2,116
Within-neighbourhood, 150 m	-0.095	0.042	0.031	2,878
Within-neighbourhood, 200 m	-0.064	0.040	0.114	3,447
Within-neighbourhood, 250 m	-0.047	0.034	0.176	3,797
Within-neighbourhood, 350 m	-0.022	0.025	0.379	4,352
Within-neighbourhood, 500 m	-0.015	0.021	0.485	5,090
Within-neighbourhood, 750 m	-0.017	0.016	0.305	4,972

Notes: Outcome is log assessed value per square metre; entries are log coefficients, not percentage transformations. Within-neighbourhood models absorb neighbourhood-by-period effects and cluster by neighbourhood, requiring near and far assessment zones in both periods. The far-control and future-site specifications instead use city effects and one-kilometre grid clustering. Area weighting uses zone area. Certain pins excludes uncertain facility locations. The 2015-site restriction removes comparison zones within 250 metres of a site opened in 2015. Radius checks use the baseline within-neighbourhood design. Original polygon multiplicity is retained. p-values use the fitted-model degrees of freedom. No multiple-radius adjustment is applied.

D.10 Reconciling the Baseline and Event-Study Specifications

Table D9: Boundary responses as cohort, distance support and reference periods change

Specification	Women (%)	Men (%)	Rows per gender
All cohorts, six bands	31.3	23.1	118,588
	[9.0, 58.0]	[7.7, 40.8]	
2017 cohort, six bands	22.9	21.9	75,987
	[-2.7, 55.2]	[3.1, 44.2]	
2017, boundary/far, pooled pre	19.2	30.6	20,722
	[-9.5, 57.0]	[7.3, 59.0]	
2017, boundary/far, 2015 reference	15.5	36.2	20,722
	[-19.3, 65.2]	[8.1, 71.5]	
2017, boundary/far: 2018	11.2	35.2	20,722
	[-28.8, 73.9]	[2.9, 77.7]	
2017, boundary/far: 2019	20.4	37.2	20,722
	[-18.2, 77.2]	[6.9, 76.0]	

Notes: Female and male expected counts are estimated by weighted PPML with cell, assigned-park-by-year and capture-month effects and park clustering. Brackets contain transformed 95 per cent confidence intervals. The first two rows use six inner bands; subsequent rows use only 0–25 and 750–1,500 metre cells. The boundary/far rows share the same 20,722 observations per gender. The six-band cohort model retains 107 additional boundary/far observations as well as the intermediate bands. Changing support is therefore part of this comparison. Pooling post periods is not an arithmetic average of the event coefficients. All female and male supports coincide within each row.

The multiple-membership boundary estimates in Table 3 are 32.9 per cent for women and 28.3

per cent for men under weights split across memberships. Giving each membership the full original weight instead gives 26.3 and 24.7 per cent; the latter male covariance requires the estimator's positive-semidefinite adjustment. These checks change the comparison area and weighting as well as membership. About 71.4 per cent of focal-post comparison pair-years are within 750 metres of another already-open park (41.8 per cent under split weights), so repeated membership does not remove interference.

E Model Data, Diagnostics and Sensitivity

I document the estimation samples, destination support and trip-making checks for Section 7, followed by facility inputs, the location-search algorithms and counterfactual sensitivities. Common layouts evaluated for both genders are reported separately from each gender's own selected layout.

E.1 Diary Construction and the Model Sample

Table E1: Travel on the model sample

	Trip-making (%)		Mean distance (km)		Trips per traveller	
	Men	Women	Men	Women	Men	Women
Work and business	53.15	21.52	5.15	4.85	1.06	1.03
School	12.42	12.12	4.35	4.12	1.01	1.01
Shopping	5.61	26.95	2.92	1.62	1.03	1.03
Broad non-work	22.86	41.50	2.73	1.88	1.17	1.21
Recreation	1.14	0.48	2.04	2.38	1.01	1.01
RPTRA composite	3.63	5.09	2.71	2.12	1.08	1.13

Notes: 2002 SITRAMP, 185,744 validated mainland Jakarta adults aged 15–110 (93,730 men and 92,014 women) before purpose-specific exclusions. Qualifying trips depart from home in the home neighbourhood and end within the 261 mainland neighbourhoods, excluding returns home. Ambiguous potentially qualifying geography removes adults for that purpose rather than coding zero; the composite excludes 20 men and 21 women, recreation three men and two women. Trip-making is the percentage of eligible adults, including non-travellers, with at least one qualifying departure. Trips per traveller condition on making such a trip; both use person weights. Mean distances use trip weights, centroid distances and equal-area-disc internal distances, not routes. Broad non-work comprises shopping, private and other; recreation covers recreation/sport/park trips; the composite adds private trips to educational and medical destinations. Purposes overlap. Point estimates only.

The roster uses complete relational identifiers. Linked and declared trip counts must agree, redundant household and person keys must be consistent, and expansion weights must be positive. The mainland adult roster contains 92,014 women and 93,730 men. The Pengadegan correction retains 389 valid adults using the source zone dictionary. Home returns and non-home-based trips remain available for complete-diary descriptions, but are excluded from model flows. Non-work legs per adult are consequently 0.648 for women and 0.463 for men in the complete diaries, compared with 0.504 and 0.268 in the model sample.

I do not code unresolved destinations as no travel. If an adult has a potentially qualifying home departure whose destination cannot be resolved, the outcome is missing for that purpose. The adult is removed from both its destination flows and its trip-making roster. The main composite excludes 41 adults: 21 women and 20 men. Their complete-diary travel status and records for other purposes are retained.

Known external destinations do not qualify. Airport, port and generic unknown zone codes are audited separately rather than automatically assigned outside the city. Table E2 counts eligible people, while Table E3 audits trip records; purpose-specific person counts overlap. Income and child-presence controls include explicit unknown categories. No child is coded only when the household roster is complete; an observed child establishes presence even if the roster is otherwise incomplete. Vehicle ownership includes sedans, vans, trucks and motorcycles. Household size uses the reported family-member count.

Table E2: Purpose-specific roster and trip frequency

Purpose	Adults F	Adults M	Excl. F	Excl. M	$\bar{m}^f$	$\bar{m}^m$
Work and business	91,880	93,056	134	674	1.031	1.060
School	91,985	93,707	29	23	1.012	1.012
Shopping	91,936	93,714	78	16	1.030	1.027
Private	91,949	93,637	65	93	1.107	1.136
Other	91,997	93,708	17	22	1.114	1.123
Broad non-work	91,857	93,600	157	130	1.215	1.172
Recreation	92,012	93,727	2	3	1.011	1.006
RPTRA composite	91,993	93,710	21	20	1.127	1.082
Extended composite	91,985	93,705	29	25	1.143	1.101

Notes: 2002 SITRAMP residents aged 15–110 in mainland Jakarta with valid complete diaries. Qualifying trips depart from home in the home neighbourhood and end within the 261 mainland neighbourhoods; home returns and known external destinations are excluded. Adults F/M count eligible women/men, including non-travellers. Excl. counts adults removed for unresolved potentially qualifying geography. $\bar{m}^g$ is the person-weighted mean qualifying legs among adults making at least one such trip. Broad non-work comprises shopping, private and other purposes. The RPTRA composite combines recreation/sport/park trips with private trips to educational or medical destinations; the extended composite adds other-purpose trips to those destinations. Purposes overlap.

Table E3: Destination geography in the audited trip records

Gender	Status	Trips	Weighted trips
Men	Known outside mainland	7,808	282,087.0
Men	Mainland mapped	211,143	7,313,724.0
Men	Unknown or unlocated	1,118	38,702.1
Women	Known outside mainland	4,067	149,063.7
Women	Mainland mapped	172,228	5,990,592.0
Women	Unknown or unlocated	415	15,099.6

Notes: Linked trips of the validated 2002 SITRAMP mainland adult roster, ages 15–110, including home returns and departures away from home. Trips counts records; weighted trips sums trip expansion factors. Destinations mapped inside the 261-neighbourhood system are distinguished from known external destinations and unresolved records. The latter are not assigned automatically to either group. This is a geography audit, not the qualifying home-departure sample. No uncertainty intervals are reported.

E.2 Destination Choice and Support

The gravity stage separates gender-specific distance decay from origin effects and common destination attractiveness. Table E4 reports both the estimates and the destination support used to construct access. A destination with no qualifying incoming flow is assigned zero attractiveness under the maintained convention. Destinations lost through separation or other support restrictions are recorded separately from zero-inflow destinations. All 261 destinations have identified positive attractiveness for the main

Table E4: Common-attractiveness gravity estimates and destination support

Purpose	δ^m	δ^f	F/M	$p_{F=M}$	Est.	Zero	Uns.
Work and business	0.346	0.373	1.08	<0.001	261	0	0
School	0.424	0.452	1.07	<0.001	261	0	0
Shopping	0.703	1.478	2.10	<0.001	261	0	0
Private	0.716	0.875	1.22	<0.001	261	0	0
Other	0.819	1.138	1.39	<0.001	260	1	0
Broad non-work	0.730	1.149	1.57	<0.001	261	0	0
Recreation	1.070	0.897	0.84	0.052	230	31	0
RPTRA composite	0.768	1.020	1.33	<0.001	261	0	0
Extended composite	0.766	1.056	1.38	<0.001	261	0	0

Notes: 2002 SITRAMP mainland residents aged 15–110 with valid complete diaries. Flows use trip expansion factors and include qualifying home departures to the 261 mainland neighbourhoods, excluding returns and known external destinations. PPML includes zero flows, gender-specific origin effects and common destination effects. δ^g is exponential decay per kilometre using centroid and equal-area-disc internal distances; F/M is female divided by male decay. Equality tests use origin-clustered covariance. Est. counts destinations with identified positive attractiveness; Zero counts destinations with no qualifying inflow assigned attractiveness zero; Uns. counts other unsupported destinations. Each grid starts with 136,242 origin–destination–gender cells. Removing zero origin-gender rows need not make a destination unsupported. Broad non-work is shopping, private and other; the RPTRA composite combines recreation with private educational/medical trips, and the extended composite adds other-purpose educational/medical trips. Purposes overlap.

Table E5: Female-to-male distance-decay ratios across specifications

Purpose	Main	Gender FE	Within .5	Within 1.5	Age 18+
Work and business	1.079	1.117	1.076	1.082	1.075
School	1.066	1.082	1.063	1.067	1.088
Shopping	2.101	2.101	1.967	2.066	2.113
Private	1.223	1.226	1.210	1.214	1.248
Other	1.390	1.403	1.350	1.379	1.458
Broad non-work	1.573	1.596	1.514	1.559	1.616
Recreation	0.838	0.836	0.851	0.845	0.921
RPTRA composite	1.329	1.351	1.305	1.318	1.399
Extended composite	1.379	1.406	1.350	1.365	1.449

Notes: Trip-weighted PPML for qualifying home departures in the 2002 SITRAMP mainland adult sample, ages 15–110. Destinations are within the 261 mainland neighbourhoods; home returns and known external trips are excluded and zero-flow cells retained. Entries divide female by male distance decay. Main uses gender-specific origin effects and common destination effects; Gender FE allows destination effects to differ by gender. Within .5 and 1.5 rescale only the equal-area-disc internal distances; cross-neighbourhood distances are unchanged. Age 18+ rebuilds the roster and flows. Median positive attractiveness is normalised to one in each fit. Broad non-work comprises shopping, private and other; the RPTRA composite combines recreation and private educational/medical trips, while the extended composite also includes other-purpose educational/medical trips. Purposes overlap. Point estimates are shown without intervals.

composite. Recreation flows are much sparser and require a separate support assessment. Zero-inflow destinations are not replaced with small positive values. Table E5 varies destination fixed effects, internal distances and the age restriction. Support in the gravity stage alone does not establish that the trip-making stage can predict for its full eligible roster.

E.3 Trip-Making, Constraints and Nonlocal Sensitivity

Table E6 reports the unrestricted access coefficient and the coefficient constrained to $[0, 1]$. When a bound binds, the other logit coefficients are re-estimated. These estimates use all purpose-eligible adults, including non-travellers. Their relation to access is conditional on the household controls and sub-district effects, not an identified causal response to new facilities. Perfect prediction removes 1,246 women from

Table E6: Unrestricted and constrained trip-making coefficients

Purpose	λ_U^f	SE F	λ_U^m	SE M	λ_C^f	λ_C^m	$p_{F=M}$
Work and business	0.119	0.063	0.011	0.076	0.119	0.011	0.198
School	-0.050	0.065	-0.059	0.074	0.000	0.000	0.903
Shopping	0.089	0.040	0.114	0.061	0.089	0.114	0.648
Private	0.269	0.115	0.227	0.141	0.269	0.227	0.611
Other	0.024	0.153	-0.312	0.148	0.024	0.000	0.001
Broad non-work	0.089	0.068	0.020	0.091	0.089	0.020	0.213
Recreation	0.348	0.282	0.197	0.187	0.348	0.197	0.593
RPTRA composite	0.317	0.084	0.108	0.094	0.317	0.108	0.015
Extended composite	0.123	0.055	0.003	0.074	0.123	0.003	0.086

Notes: Person-weighted logits for at least one qualifying home departure in the 2002 SITRAMP mainland sample, ages 15–110, including eligible non-travellers. Destinations must be within the 261 neighbourhoods; home returns and known external trips are excluded. λ_U is the unrestricted log-access coefficient. λ_C restricts it to $[0, 1]$, refitting the other coefficients when a bound binds. Controls are income band, vehicle ownership, household size, children under fifteen with unknown categories, and sub-district effects. Standard errors and pooled female–male equality tests cluster by origin and condition on fitted access. The pooled test interacts controls and fixed effects with gender. Recreation drops 1,246 women through perfect prediction and is diagnostic. Broad non-work comprises shopping, private and other; the RPTRA composite is recreation plus private educational/medical trips, and the extended composite adds other-purpose educational/medical trips. Purposes overlap.

the recreation logit, or 1,161 in the age-18+ check. I do not fill in predictions for those excluded adults. The three purposes used in counterfactuals have no such exclusions in the original sample. For the main composite, fitted neighbourhood trip-making has weighted correlations of 0.477 for women and 0.512 for men, and RMSEs of 1.94 and 1.61 percentage points. These are in-sample fit diagnostics, not external validation.

The nonlocal check removes own-neighbourhood and adjacent-neighbourhood pairs from gravity estimation rather than setting their flows to zero. Thirty-one composite destinations then lose all incoming support. The hybrid aligns the remaining nonlocal destination effects with the baseline on shared support and retains baseline attractiveness for those 31 destinations. They account for 5.91 per cent of weighted female access and 5.83 per cent of weighted male access. A fully nonlocal citywide composite comparison is therefore unsupported. The hybrid tests sensitivity to nonlocal information while retaining some local information.

Table E7: Age and nonlocal trip-making sensitivities

Specification	λ_U^f	λ_U^m	F–M	$p_{F=M}$
Composite, main	0.317	0.108	0.209	0.015
Composite, age 18+	0.343	0.157	0.186	0.041
Composite, hybrid	0.331	0.058	0.273	0.005
Broad non-work, main	0.089	0.020	0.069	0.213
Broad non-work, nonlocal	-0.018	-0.034	0.016	0.849

Notes: Unrestricted log-access coefficients in person-weighted trip-making logits for the 2002 SITRAMP eligible mainland roster, ages 15–110. Qualifying home departures end within the 261 neighbourhoods and exclude returns and known external destinations. F–M is the female-minus-male coefficient. Household income, vehicles, size, child-presence/unknown categories and sub-district effects are included; equality tests cluster by origin and condition on fitted access. Age 18+ restricts both stages. Nonlocal gravity omits own and adjacent pairs. The composite hybrid aligns supported nonlocal destination effects with the baseline and retains baseline attractiveness for 31 destinations losing support. Broad non-work has supported nonlocal access, but both unrestricted coefficients become negative and the constrained coefficients are zero. Purpose definitions are broad non-work (shopping/private/other), the RPTRA composite (recreation plus private educational/medical trips), and the extended composite (also other-purpose educational/medical trips); they overlap.

E.4 Facility Inputs and Location Search

The facility inputs retain all 315 sites. Areas come from polygons for 298 sites and from the register for 17, with the latter accounting for 3.83 per cent of total area. Ceria is recorded as 1,261 m^2, correcting the interpretation of the register value 1.261. Sixteen polygon/register comparisons differ by more than a factor of two; the polygon takes precedence. Median area is 1,559 m^2, total area is 557,796 m^2, and the observed range is approximately 80–6,763 m^2.

Validated pins determine host neighbourhoods, and the baseline represents each site by its host centroid. The pin-to-centroid shift averages 753 metres, with a median of 655 metres. The 315 sites occupy 168 neighbourhoods. Concentrating the same land in one park gives 55.78 hectares; three parks have 18.59 hectares each and six have 9.30 hectares each. All are outside the observed RPTRA size range. Hypothetical large parks occupy distinct candidate centroids and retain area $\bar{S}/K$ throughout the search.

At $\rho = 1$, equal area gives equal total attractiveness. Other size elasticities change both total attractiveness and its distribution. The central location is defined using distance to Monas. A one-park search evaluates all 261 candidate centroids. Three- and six-park searches use greedy addition and one pass of single-site swaps, with deterministic tie-breaking. The search returns best-found layouts rather than proven global optima. Scores are screened numerically and close candidates evaluated exactly. The search imposes no land-feasibility, cost, congestion or capacity constraints. Selected locations vary across

Table E8: Preferred candidate neighbourhoods at the main calibration

Target	K	Neighbourhoods
Men	1	Ciracas
Men	3	Kelapa Dua Wetan, Penggilingan, Tugu Utara
Men	6	Kelapa Dua Wetan, Ciracas, Susukan, Penggilingan, Tugu Utara, Semper Barat
Women	1	Tugu Utara
Women	3	Kelapa Dua Wetan, Penggilingan, Tugu Utara
Women	6	Kelapa Dua Wetan, Susukan, Penggilingan, Cakung Barat, Tugu Utara, Lagoa

Notes: Counterfactual predictions combine 2002 SITRAMP estimates, fixed 2010 adult populations and 315 registered RPTRA sites. The purpose is recreation plus private educational/medical trips. Networks use 55.78 hectares and 261 candidate centroids, relative to the destination set with no added network. The calibration is $\alpha = .05$, $\rho = 1$: a median-area facility has five per cent of median positive destination attractiveness, with unit area elasticity. K is the number of equal-area large parks. Each row selects locations to maximise its target gender's net additional within-purpose trips. The one-park search is exhaustive; three and six parks use greedy addition and one swap pass. Locations are point-estimate choices, not verified feasible parcels or uniquely stable recommendations. No location intervals are reported.

bootstrap draws. For women, the one-park choice is Tugu Utara in 44.0 per cent of draws, Penggilingan in 29.0 per cent and Ciracas in 23.5 per cent. For men, conditional on a positive trip-making coefficient, Ciracas is selected in 58.0 per cent and Tugu Utara in 25.6 per cent. Zero-coefficient draws leave all locations tied and are excluded from these frequencies. The frequencies describe uncertainty within the candidate-centroid model rather than a uniquely preferable parcel.

E.5 Attractiveness and Purpose Mapping

Table E9 varies the attractiveness calibration and the set of activities represented by a facility. Because the purpose definitions overlap, I estimate and report them separately. Table E10 reports every detected crossing between the observed network and one central park over the stated size-elasticity grid. The main composite retains the same equal-attractiveness ranking across the three attractiveness

Table E9: Equal-pull outcomes across attractiveness and purpose definitions

Purpose	α	Obs. F	Cen. F	Best F	Obs. M	Cen. M	Best M
RPTRA composite	0.01	1.05	0.22	1.98	0.24	0.07	0.51
RPTRA composite	0.05	4.99	1.07	6.60	1.15	0.34	1.60
RPTRA composite	0.10	9.48	2.03	10.19	2.18	0.65	2.44
Broad non-work	0.01	1.56	0.34	2.34	0.23	0.09	0.38
Broad non-work	0.05	7.34	1.58	7.27	1.10	0.43	1.29
Broad non-work	0.10	13.76	2.94	10.73	2.09	0.81	1.97
Extended composite	0.01	0.60	0.11	1.06	0.01	0.00	0.02
Extended composite	0.05	2.81	0.53	3.01	0.04	0.01	0.05
Extended composite	0.10	5.26	0.99	4.44	0.07	0.02	0.07

Notes: Model predictions use 2002 SITRAMP estimates, fixed 2010 populations and the 315-site register. Every network uses 55.78 hectares and 261 candidate centroids. Gains are relative to no added network. Entries are net additional trips within the stated purpose, in thousands per day, excluding trips diverted from existing destinations; F/M denote women/men. α sets a median-area facility's attractiveness relative to median positive destination attractiveness, and $\rho = 1$ makes attractiveness proportional to area. Obs. retains observed sites and areas at host centroids; Cen. is one park at the centroid nearest Monas; Best is one park at the target gender's preferred location. Broad non-work is shopping/private/other. The RPTRA composite combines recreation with private educational/medical trips; the extended composite adds other-purpose educational/medical trips. Purposes overlap. Point predictions have no intervals.

Table E10: All original-sample crossings against one central park

Purpose	Gender	$\alpha = .01$	$\alpha = .05$	$\alpha = .10$
RPTRA composite	Women	1.290	1.317	1.338
RPTRA composite	Men	1.226	1.244	1.260
Broad non-work	Women	1.290	1.334	1.372
Broad non-work	Men	1.173	1.189	1.204
Extended composite	Women	1.319	1.358	1.388
Extended composite	Men	1.232	1.252	1.269

Notes: Counterfactuals combine 2002 travel estimates, fixed 2010 adult populations and the 315-site register on 261 candidate centroids, holding land at 55.78 hectares. α gives median-area facility attractiveness relative to median positive destination attractiveness. Entries are size elasticities ρ at which observed-network and one-central-park net additional trips are equal for the stated purpose and gender. Central means the centroid nearest Monas. The search scans $[.5, 2]$ in steps of .05 and refines sign changes by bisection. All detected crossings favour observed sites below and the central park above; tangencies or closely spaced double crossings can be missed. Broad non-work is shopping/private/other; the RPTRA composite is recreation plus private educational/medical trips; the extended composite adds other-purpose educational/medical trips. These overlapping-purpose point scenarios have no intervals.

calibrations. The ordering changes for some other purposes and comparison networks. Broad non-work trip-making is particularly uncertain, and its bootstrap interval for the gender difference includes zero. The extended composite's unrestricted female-minus-male difference is 0.120, with interval [0.001, 0.251]. Recreation has a different purpose definition and lacks supported full-roster trip-making predictions. These comparisons do not estimate the effect of replacing recreation with a different facility bundle.

E.6 Size Sensitivity

Table E11 fixes attractiveness calibration at $\alpha = .05$ and varies the area elasticity. It compares observed site areas with one central park, without optimising the large park's location. Equal land produces equal total attractiveness only at $\rho = 1$, so other rows vary returns to size as well as spatial concentration.

Table E11: Full size grid for the main central-park comparison

ρ	Observed F	Central F	Observed M	Central M
0.50	4.39	0.06	1.02	0.02
0.55	4.43	0.08	1.03	0.03
0.60	4.48	0.11	1.04	0.03
0.65	4.53	0.14	1.05	0.05
0.70	4.58	0.19	1.06	0.06
0.75	4.64	0.26	1.07	0.08
0.80	4.70	0.34	1.09	0.11
0.85	4.77	0.46	1.10	0.15
0.90	4.83	0.61	1.12	0.20
0.95	4.91	0.81	1.13	0.26
1.00	4.99	1.07	1.15	0.34
1.05	5.07	1.40	1.17	0.45
1.10	5.16	1.84	1.19	0.59
1.15	5.25	2.40	1.21	0.77
1.20	5.34	3.12	1.23	1.00
1.25	5.45	4.02	1.25	1.29
1.30	5.55	5.15	1.28	1.65
1.35	5.67	6.56	1.30	2.09
1.40	5.79	8.28	1.33	2.62
1.45	5.91	10.37	1.36	3.26
1.50	6.04	12.89	1.39	4.02
1.55	6.18	15.91	1.42	4.92
1.60	6.33	19.47	1.45	5.95
1.65	6.48	23.65	1.49	7.15
1.70	6.64	28.51	1.52	8.51
1.75	6.81	34.12	1.56	10.04
1.80	6.98	40.55	1.60	11.77
1.85	7.17	47.87	1.64	13.68
1.90	7.36	56.14	1.69	15.80
1.95	7.56	65.46	1.73	18.14
2.00	7.77	75.87	1.78	20.68

Notes: 2002 SITRAMP estimates, fixed 2010 populations and the RPTRA register determine predictions for recreation plus private educational/medical trips. All layouts use 55.78 hectares and 261 candidate centroids, relative to no added network. Entries are net additional within-purpose trips in thousands per day; F/M denote women/men. A median-area site has attractiveness $\alpha = .05$ relative to median positive destination attractiveness; the rows vary its area elasticity ρ. Observed uses original site areas at host centroids; Central places all land at the centroid nearest Monas. Total attractiveness is equal only at $\rho = 1$. These are point predictions, without intervals or optimisation of the central location.

E.7 Common Layouts and Counterfactual Uncertainty

I calculate the female share of gains on common networks. Table E12 evaluates each gender-targeted layout for both genders; Table E13 reports uncertainty for each gender's own selected layout. Using

Table E12: Both genders evaluated on each selected network

Optimisation target	K	Trips F	Trips M
Men	1	5.969	1.598
Men	3	8.627	1.953
Men	6	8.441	1.992
Women	1	6.596	1.422
Women	3	8.627	1.953
Women	6	8.662	1.944

Notes: Predictions combine 2002 travel estimates, fixed 2010 adult populations and the RPTRA register for recreation plus private educational/medical trips. Each network uses 55.78 hectares on 261 candidate centroids, relative to no added network. The calibration is $\alpha = .05$, $\rho = 1$, so attractiveness is proportional to area and a median-area facility has five per cent of median positive destination attractiveness. Trips F/M are net additional female/male trips in thousands per day; K is the number of equal-area large parks. Each layout is selected for the named target and then evaluated for both genders. One-site searches are exhaustive; three- and six-site searches use greedy addition and single swaps. Each row therefore describes one common layout. Point estimates have no intervals.

Table E13: Sampling uncertainty for the main layouts

Layout	Point F	95% interval F	Point M	95% interval M
315 small, observed	4.99	[2.53, 7.63]	1.15	[0.00, 3.11]
1 large, central	1.07	[0.51, 1.79]	0.34	[0.00, 0.97]
1 large, best location	6.60	[3.24, 12.49]	1.60	[0.00, 4.90]
3 large, best found	8.63	[4.10, 14.26]	1.95	[0.00, 5.66]
6 large, best found	8.66	[4.19, 14.18]	1.99	[0.00, 5.67]

Notes: Model predictions for recreation plus private educational/medical trips use 2002 estimates, fixed 2010 populations and 55.78 hectares from the RPTRA register, on 261 candidate centroids. Gains are relative to no added network. The calibration is $\alpha = .05$, $\rho = 1$, giving equal total attractiveness. Entries are net additional trips in thousands per day, separately for women and men. Central is the centroid nearest Monas; best rows select each gender's own layout, exactly for one park and heuristically for three or six. Intervals are the 2.5th and 97.5th percentiles of 200 origin-neighbourhood bootstrap draws within sub-districts. Both genders are resampled jointly, both stages re-estimated and locations reselected in every draw. Population and calibration inputs stay fixed; zero-response draws are retained.

matched draws, the female observed-minus-central gap is 3,920 trips per day [1,820, 5,901]. The female observed-minus-best-single-park gap is $-1,610$ [$-5,261$, -680]. The observed-minus-central advantage is 3,113 trips larger for women than for men [885, 4,723]. The central comparison holds the layouts common across genders; the best-single-park comparison selects each gender's own location. Three- and six-park point estimates are close, and these results do not establish a preference for three over six.

E.8 Any-Travel Illustration and Street-Occupancy Comparison

An additional trip-maker for the modelled purpose may otherwise have travelled for another purpose or made no trip. Let H_i^g be the observed no-travel share for neighbourhood i and gender g. Holding trip-making in other purposes fixed, an illustrative allocation envelope is

$$0 \;\le\; \Delta V^g \;\le\; \sum_i N_i^g \min\{\Delta P_i^g, H_i^g\}.$$

Here ΔV^g counts additional adults making any trip. The proportional illustration is $\sum_i N_i^g \bar{h}_i^g \Delta P_i^g$, where $\bar{h}_i^g$ is the no-travel share among adults not making a trip for the purpose. Cells with fewer than 30 such adults use a sub-district pool, and the implementation respects the available non-traveller caps. This is an assumed allocation of additional purpose-specific trip-makers, not an identified change in women leaving home. It does not enter the planning objective. Trips added within the modelled purpose may substitute for work, shopping or other journeys. I do not estimate cross-purpose substitution or convert arrivals into visible street time, so agreement in aggregate ratios across sources cannot identify either margin.

Table E14: Illustrative any-travel envelope for the observed network

Gender	Purpose change	Lower	Upper	Proportional
Women	0.1273	0	0.1273	0.0329
Men	0.0302	0	0.0302	0.0042

Notes: The observed RPTRA network is added to the 2002 destination set using the composite purpose, fixed 2010 adult populations, 55.78 hectares and 261 candidate centroids. The calibration is $\alpha = .05$, $\rho = 1$. Entries are percentage points of the relevant adult population. Purpose change is the predicted increase in adults making a composite trip. Holding other-purpose trip-making fixed, Lower is zero and Upper sums origin-level increases capped by the observed no-travel share. Proportional uses the no-travel share among purpose non-trip-makers, retaining that cap and pooling cells with fewer than 30 purpose non-trip-makers by sub-district. These are illustrative allocations, not identified changes in people leaving home. No bootstrap intervals are reported.

Table E15: Sensitivity of the street-presence consistency comparison

α	Addition	Change F (%)	Change M (%)	F/M
0.01	All sites together	0.86	0.58	1.50
0.01	One site at a time	0.92	0.93	0.99
0.05	All sites together	4.12	2.72	1.52
0.05	One site at a time	4.59	4.64	0.99
0.10	All sites together	7.85	5.09	1.54
0.10	One site at a time	9.12	9.23	0.99

Notes: Modelled changes in arrivals to RPTRA host neighbourhoods use 2002 estimates for recreation plus private educational/medical trips and fixed 2010 populations. Each addition is a median-area unit of 1,559 square metres at an observed site, with $\rho = 1$ and attractiveness α relative to median positive destination attractiveness. All-sites rows add a unit at each of 315 sites simultaneously and divide the change in host arrivals by baseline arrivals to those hosts. One-at-a-time rows add one unit per experiment and aggregate with site-baseline-flow weights, counting repeated hosts separately. F/M divides the female by the male percentage change. These point scenarios have no intervals; they predict neither street counts nor the fixed-total-area comparison.

E.9 Supporting Evidence on Site Size

I estimate heterogeneity in the per-cell boundary response by facility area. This outcome differs from total visits and destination attractiveness, leaving the model's area elasticity unidentified. The

Table E16: Predicted per-cell opening responses by site area

Gender	Area percentile	Area (m^2)	Response (%)	95% interval
Male	P25	1,015	25.5	[5.8, 48.8]
Male	P75	2,191	19.2	[1.7, 39.9]
Female	P25	1,015	42.0	[11.2, 81.2]
Female	P75	2,191	18.0	[-3.0, 43.6]

Notes: Jakarta Google Street View expected adult pedestrian counts in 2013, 2015, 2018 and 2019. Weighted PPML includes cell, nearest-site-by-year and capture-month fixed effects, sampling and observation weights and site-clustered standard errors. Post-opening boundary-band interactions compare 0–25 metres with 750–1,500 metres using the corrected Ceria geometry. The specification interacts treatment with centred log site area, log 2014 neighbourhood population density and distance to Monas. P25/P75 are the 25th/75th percentiles of site area; density and centrality are held at their all-site means. Response is $100[\exp(\widehat{\beta}) - 1]$ per cell, with transformed 95 per cent intervals. Both gender fits use 118,588 observations across 314 retained sites. Size is not randomly assigned, and these responses do not estimate the model's attractiveness elasticity.

boundary regressions retain cell, site-by-year and month effects, sampling and observation weights and site clustering. They add post-opening band interactions with centred log area, log 2014 neighbourhood density and distance to Monas. All 315 sites enter the input panel; 314 remain in estimation and 208 contribute retained observations within 25 metres. Both gender fits use 118,588 observations. Sirih Kuning's eight input observations are excluded by PPML, as in the original baseline. The Ceria correction changes the geometry on the existing cell frame.

The log-area interaction is -0.240 for women ($p = 0.069$) and -0.067 for men ($p = 0.456$). Their difference is -0.173 ($p = 0.102$). The coefficients describe variation in the per-cell response with area, conditional on density and centrality controls. Site size is not randomly assigned.

Table E17: Area-tercile sensitivity

Gender	Tercile	Response (%)	95% interval
Male	Small	8.5	[-15.1, 38.7]
Male	Medium	57.2	[17.6, 110.1]
Male	Large	10.7	[-11.2, 38.0]
Female	Small	33.3	[-5.3, 87.6]
Female	Medium	74.6	[21.8, 150.4]
Female	Large	-2.7	[-26.8, 29.4]

Notes: Jakarta Google Street View expected adult pedestrian counts in 2013, 2015, 2018 and 2019. Weighted PPML includes cell, nearest-site-by-year and capture-month fixed effects, sampling and observation weights and site-clustered standard errors. Post-opening boundary-band interactions compare 0–25 metres with 750–1,500 metres using the corrected Ceria geometry. Site-area terciles replace the continuous size interaction; cutoffs are 1,142 and 1,883 square metres. Responses are $100[\exp(\widehat{\beta}) - 1]$ for each tercile with transformed 95 per cent intervals. The specification retains the density and centrality adjustments. Female joint equality has $p = .040$ and largest-versus-smallest $p = .148$. These are per-cell street responses, not total visits or estimates of the model size elasticity.

E.10 Bootstrap Implementation and Diagnostics

The bootstrap draws 261 origins with replacement within 42 sub-districts, jointly for women and men, using seed 2026100907. Repeated origins receive distinct estimation identifiers but retain their geography. Each draw re-estimates gravity slopes, common destination attractiveness, access, trip-making and trips per trip-maker. Evaluation uses the original eligible roster and fixed 2010 populations. Exact aggregation of identical household-control/outcome records preserves the weighted logit likelihood; the implementation also checks predictions and counterfactuals against direct individual-record evaluation.

All 600 purpose-draws complete successfully, with none replaced. Destinations receiving no incoming flow in a draw are assigned zero attractiveness under the maintained convention; up to eight become newly zero. Other losses of support count as failures. Access coefficients outside $[0, 1]$ are refitted at the bound. For the illustrative no-traveller allocation only, an unsampled origin uses its sampled sub-district pool. That pooling does not affect optimisation. One-site searches remain exhaustive, while three- and six-site searches remain heuristic.

Table E18: Bootstrap fit and constraint diagnostics

Purpose	Gender	Successful	At zero	At one
RPTRA composite	Men	200	24	0
RPTRA composite	Women	200	0	0
Extended composite	Men	200	87	0
Extended composite	Women	200	0	0
Broad non-work	Men	200	77	0
Broad non-work	Women	200	19	0

Notes: Each entry counts outcomes among 200 attempted origin-neighbourhood bootstrap draws within sub-districts, jointly resampling both genders in the 2002 SITRAMP mainland sample, ages 15–110. Both estimation stages are refitted for qualifying home departures to the 261 neighbourhoods, retaining eligible non-travellers and excluding returns/external destinations. Successful counts completed fits. At zero/one counts log-access coefficients constrained to a bound in the person-weighted trip-making logit; a zero response is a valid estimate, not a failed draw. Controls cover income, vehicle ownership, household size, child presence and unknown categories, plus sub-district effects. Broad non-work is shopping/private/other; the RPTRA composite is recreation plus private educational/medical trips; the extended composite adds other-purpose educational/medical trips. Purposes overlap.

Table E19: Main-composite thresholds and conditional intervals

Gender	α	Threshold	95% interval	Crossing	Flat draws
Women	0.01	1.290	[1.234, 1.322]	100%	0
Women	0.05	1.317	[1.257, 1.347]	100%	0
Women	0.10	1.338	[1.273, 1.369]	100%	0
Men	0.01	1.226	[1.179, 1.270]	88%	24
Men	0.05	1.244	[1.194, 1.287]	88%	24
Men	0.10	1.260	[1.208, 1.300]	88%	24

Notes: Composite-purpose predictions combine 2002 travel estimates, fixed 2010 populations and 55.78 hectares from the RPTRA register on 261 candidate centroids. α calibrates median-area facility attractiveness relative to median positive destination attractiveness. Threshold is the size elasticity equating observed-network and one-central-park net additional trips; the central centroid is nearest Monas and lower elasticities favour observed sites. Intervals are 2.5th–97.5th percentiles of 200 origin-neighbourhood bootstrap draws within sub-districts, jointly resampling genders and refitting destination choice and trip-making. Populations and calibration inputs stay fixed. Crossing is the percentage of draws with a crossing in $[.5, 2]$; Flat means zero constrained trip-making response for all layouts. Intervals condition on a crossing: all 200 female draws and 176 male draws. No other no-crossing or multiple-crossing case occurs for this purpose.

For the extended composite, crossings occur in 100 per cent of female and 56.5 per cent of male draws. For broad non-work, the frequencies are 90.5 and 61.5 per cent, at each attractiveness calibration.

Remaining draws have zero trip-making responses rather than an additional nonzero-response case with no crossing in range. These are frequencies among bootstrap draws, not posterior probabilities that one policy is better. The intervals condition on geography, population, measured areas, purpose definitions and calibration. They do not resolve causal identification or uncertainty about actual facility attractiveness and returns to size.

E.11 Optimised Small Networks and Best-Location Thresholds

Network comparisons use the same 200 parameter draws. Each draw reselects layouts and evaluates them on the original roster and fixed 2010 populations. The nine-unit cap comes from the observed maximum; it is a design restriction, not an estimated physical capacity. The common population allocation is uncapped and uses total adults of both genders, with ties resolved by neighbourhood code. Greedy addition followed by one sequential swap pass need not find the global optimum.

Table E20: Break-even size elasticity with locations reselected

Gender	α	Observed / best 1	Capped small / best 1	Capped small / best 3
Women	0.01	0.864	0.987	0.955
Women	0.05	0.925	1.054	0.980
Women	0.10	0.978	1.104	1.001
Men	0.01	0.838	0.968	0.957
Men	0.05	0.911	1.043	0.987
Men	0.10	0.965	1.088	1.001

Notes: Composite-purpose counterfactuals use 2002 estimates, fixed 2010 populations and 55.78 hectares from the RPTRA register on 261 candidate centroids. α calibrates median-area facility attractiveness relative to median positive attractiveness. Entries are break-even area elasticities for net additional trips. Observed retains the existing network; Capped small reallocates 315 equal units with at most nine per neighbourhood; Best 1/3 allocates the land to one/three equal parks. Optimised layouts are reselected at every elasticity, separately by gender. The search scans $[.5, 2]$ in steps of .05 and refines sign changes to 10^{-7}. Detected crossings favour small below and large above. Multi-site searches are heuristic and the grid may miss crossings. These thresholds have no estimated bootstrap intervals.

Table E21: Matched bootstrap differences: optimised small minus large

Gender	Small design	Large parks	Gap (000s)	95% interval
Men	Cap 9	1	0.27	[-0.09, 0.63]
Men	Cap 9	3	-0.09	[-0.55, 0.00]
Women	Cap 9	1	1.45	[-0.45, 1.68]
Women	Cap 9	3	-0.58	[-1.70, -0.02]
Men	Uncapped	1	0.42	[0.00, 1.08]
Men	Uncapped	3	0.07	[0.00, 0.16]
Women	Uncapped	1	2.24	[0.88, 2.91]
Women	Uncapped	3	0.21	[0.11, 0.57]

Notes: Composite-purpose predictions combine 2002 estimates, fixed 2010 populations and 55.78 hectares from the RPTRA register on 261 centroids. The calibration is $\alpha = .05$, $\rho = 1$, giving equal total attractiveness. Gaps are small-minus-large net additional trips in thousands per day. Small layouts reallocate 315 equal units with a nine-unit cap or without a co-location limit; large layouts contain one or three equal parks. Each bootstrap draw reselects both layouts for the indicated gender. Intervals use paired differences and the 2.5th–97.5th percentiles of 200 origin-neighbourhood draws within sub-districts, resampling genders jointly and refitting both stages. Populations and calibration inputs stay fixed. Zero-response draws are retained. The intervals are not differences between marginal interval endpoints.

E.12 Educational and Medical Component

Private trips to educational and medical destinations have distance decay of 1.038 for women and 0.694 for men, a ratio of 1.50. Their access coefficients are 0.333 and 0.393, with $p = 0.490$ for conditional equality. This component has no analogous female advantage in the access coefficient. The codes do not identify escorting or care, and coefficients from separately estimated purpose outcomes do not add up to the composite coefficient.

Table E22: Private trips to educational and medical destinations

Gender	Eligible	Retained	Access coefficient	Std. error
Men	93,713	93,713	0.393	0.100
Women	91,995	91,995	0.333	0.077

Notes: 2002 SITRAMP mainland adults aged 15–110 with valid complete diaries. The outcome is at least one home departure for private purpose (code 5) to an educational or medical destination (facility codes 4 or 6) within the 261 mainland neighbourhoods. Returns and known external trips are excluded. Eligible counts the purpose-specific roster, including non-travellers; Retained counts adults entering the person-weighted logit. Entries are log-access coefficients with origin-clustered errors conditional on fitted access. Gravity uses trip weights, gender-specific origins and common destination effects. Logit controls include income, vehicles, household size, child presence and unknown categories, plus sub-district effects. No component bootstrap is reported. The diary codes do not directly identify escorting or caregiving.

E.13 Geographical Distribution and Actual-Pin Sensitivity

Figures E2 and E3 in the main text map predicted gains. Table E23 replaces host-centroid locations with actual site pins when calculating distances from other neighbourhoods. Host-neighbourhood distance retains the baseline convention. The sensitivity varies the location of observed facilities, holding baseline access and behavioural coefficients fixed.

Table E23: Actual-pin sensitivity for the observed network

Gender	Site representation	Trips (000s)	Baseline %
Men	centroids	1.151	0.834
Men	actual pins	1.184	0.858
Women	centroids	4.986	2.512
Women	actual pins	5.238	2.639

Notes: The observed-network sensitivity uses 2002 composite-purpose estimates, fixed 2010 populations and the 315 RPTRA sites totalling 55.78 hectares, relative to no added network. The calibration is $\alpha = .05$, $\rho = 1$. Trips are net additional within-purpose trips in thousands per day. Baseline percentages divide by 198,482 female and 137,957 male modelled trips. Centroids places each facility at its host neighbourhood centroid. Actual pins uses origin-centroid to site-pin great-circle distances for non-host origins; host origins retain the equal-area-disc distance. Baseline access and behavioural coefficients are unchanged. The geometry sensitivity has no bootstrap intervals.

E.14 Distance Sensitivity across Trip Purposes

Writing $\eta_{jq} = \ln A_{jq}$, the reduced-form gravity notation is

$$\mathbb{E}\left[T_{ijq}^{g} \mid \alpha_{iq}^{g}, \eta_{jq}, d_{ij}\right] = \exp\left(\alpha_{iq}^{g} + \eta_{jq} - \delta_{q}^{g} d_{ij}\right), \tag{E1}$$

Origin effects and distance slopes vary by gender, and the destination effects are common within purpose. This is the same specification as Equation (7).

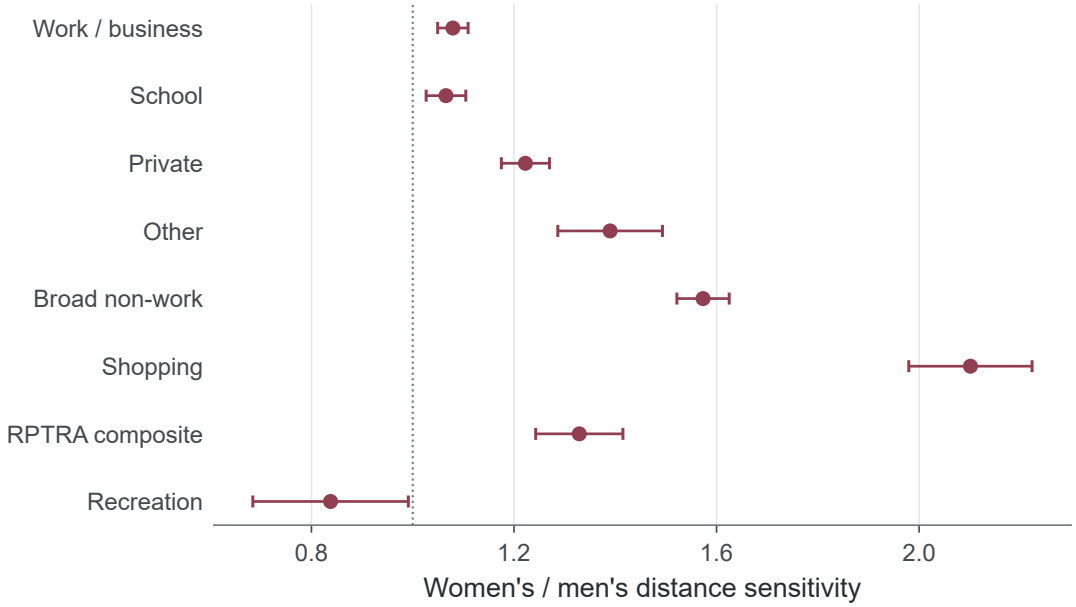
Figure E1 reports the conditional distance-decay comparisons introduced in Section 7.2. The main ratios use gender-specific origins and common destination effects; Table E5 varies destination attractiveness, internal distance and the adult age restriction. The broad non-work ratio is 1.57, compared with 1.08 for work/business and 1.07 for school. Recreation alone has a ratio of 0.84. The gradient is therefore purpose-specific rather than a common female travel penalty.

E.15 Comparisons across Facility Numbers and Sizes

These comparisons vary facility number and size as well as location, unlike the fixed-100-site comparison in Section 7.6. They retain the earlier layout definitions and calibration exercises. All reported results remain conditional on the same historical travel parameters, fixed population and specified attractiveness assumptions.

Layouts and the land constraint. Total area is $\bar{S} = 557{,}796$ square metres, or 55.78 hectares, the combined area of the 315 registered facilities. The median site area is $s_0 = 1{,}559$ square metres. The observed network retains each facility's area and host neighbourhood, placing it at that neighbourhood's centroid. It occupies 168 neighbourhoods. Every scenario uses the same 261 candidate centroids and the same internal-distance convention.

Figure E1: Gender differences in destination distance decay, by trip purpose



Notes: 2002 SITRAMP valid-diary residents aged 15–110, with trips departing from home and ending within Jakarta's 261 mainland neighbourhoods, using trip expansion weights. Points divide the female distance slope by the male slope in PPML with gender-specific origin effects and common destination effects; zero flows are included. Distances are in kilometres. Horizontal bars are 95 per cent delta-method intervals using covariance clustered by origin. These differ from the joint two-stage bootstrap intervals used for selected model quantities. Broad non-work is shopping, private and other trips; the RPTRA composite combines recreation with private trips to educational and medical destinations. The extended composite also includes other-purpose trips to educational and medical destinations. Purposes overlap; return-home legs are excluded.

Large-park scenarios divide the land equally among one, three or six parks. Locations are either central, measured relative to Monas, or selected to maximise additional trips for a given gender. The one-park search evaluates every candidate. Multi-park searches add sites sequentially and then test single swaps, so they identify the best layouts found rather than verified global optima. Large parks occupy distinct neighbourhoods. Locations are selected separately for women and men, and each resulting network is evaluated for both.

The small-unit scenarios reallocate 315 equal facilities of area $\bar{S}/315$, approximately $1{,}771$ square metres each. The main search permits at most nine units per neighbourhood, the observed maximum, and uses greedy addition followed by one pass of single-unit swaps. An uncapped sensitivity permits further co-location. A population benchmark allocates units in proportion to total adult population using largest remainders. All scenarios use the same total land but impose different restrictions on its distribution.

Facility size and spatial reach. Two considerations affect the choice between concentrated and dispersed facilities. Returns to size determine their total attractiveness. Its location then determines who can reach it and how much additional access each origin gains. Concentration can produce a large improvement in a few neighbourhoods but leave others far away. Since trip-making is concave in access under the imposed coefficient range, very large access gains at one origin have diminishing effects on its trip-making probability.

I summarise the size comparison using the break-even elasticity ρ^{*g} at which two layouts generate the same net additional trips for gender g. This is a threshold conditional on α, not an estimate of actual returns to area. For each comparison, I report every crossing in the scanned range and which layout

performs better on either side. Best-location layouts are reselected as ρ changes. A crossing need not exist or be unique. When $\lambda_q^g = 0$, no network changes trip-making for that gender, so no unique trip-making threshold exists. The boundary-size regressions in Appendix E.9 cannot identify ρ: a change in counts per street cell is not a change in site attractiveness.

Table E24: Net additional trips under alternative allocations of the same public-space land

Network	Trips (000s)	95% interval	Baseline %	Trips/site
Women				
315 small, observed	4.99	[2.53, 7.63]	2.51	15.8
1 large, central	1.07	[0.51, 1.79]	0.54	1,066.0
1 large, best location	6.60	[3.24, 12.49]	3.32	6,596.1
3 large, best found	8.63	[4.10, 14.26]	4.35	2,875.7
6 large, best found	8.66	[4.19, 14.18]	4.36	1,443.6
315 small, best found (cap 9)	8.05	[3.98, 12.77]	4.06	25.6
315 small, best found (uncapped)	8.83	[4.29, 14.71]	4.45	28.0
315 small, population allocation	4.62	[2.36, 6.96]	2.33	14.7
Men				
315 small, observed	1.15	[0.00, 3.11]	0.83	3.7
1 large, central	0.34	[0.00, 0.97]	0.25	343.6
1 large, best location	1.60	[0.00, 4.90]	1.16	1,597.8
3 large, best found	1.95	[0.00, 5.66]	1.42	651.1
6 large, best found	1.99	[0.00, 5.67]	1.44	332.0
315 small, best found (cap 9)	1.87	[0.00, 5.25]	1.35	5.9
315 small, best found (uncapped)	2.02	[0.00, 5.78]	1.47	6.4
315 small, population allocation	1.08	[0.00, 2.85]	0.78	3.4

Notes: Model predictions using 2002 SITRAMP travel estimates, 2010 adult populations and the register of 315 RPTRA facilities. The purpose combines recreation and private trips to educational and medical destinations. Every network uses 55.78 hectares. Facility attractiveness is 5 per cent of median positive destination attractiveness ($\alpha = .05$), and the size elasticity is one ($\rho = 1$). Trips (000s) are thousands of net additional trips within this purpose per day; trips displaced from existing destinations are excluded. Baseline % divides by baseline modelled trips: 198,482 for women and 137,957 for men. Trips/site divides by the number of facilities, not occupied neighbourhoods. Intervals are 2.5th–97.5th percentiles of 200 origin-neighbourhood bootstrap draws within sub-districts, resampling both genders jointly, re-estimating both model stages and reselecting locations; population and calibration inputs remain fixed. Gains are relative to no added network; central is the candidate centroid nearest Monas. Best rows use each gender's own layout. Small capped and uncapped rows occupy 38/9 neighbourhoods for women and 37/8 for men. One-large search is exact over 261 candidates; all multi-unit searches are best found. Cap 9 allows at most nine equal units per neighbourhood; population allocation uses total adult population proportions. The uncapped row permits unequal total area across neighbourhoods through co-located units.

The observed network. At $\alpha = 0.05$ and $\rho = 1$, the observed network generates 4.99 thousand net additional composite trips per day for women and 1.15 thousand for men (Table E24). The respective 95 per cent bootstrap intervals are $[2.53, 7.63]$ and $[0, 3.11]$ thousand. These gains equal 2.51 and 0.83 per cent of baseline composite trips, or 15.8 and 3.7 additional trips per facility per day.

Most predicted visits to the new facilities are diverted from existing destinations. Net additional trips account for 31.0 per cent of women's trips to the new network and 11.0 per cent of men's. The female advantage in additional trip-making therefore differs from the gender composition of total facility visits. Neither quantity is inferred from the interior photographs.

Alternative locations. One central park generates 1.07 thousand additional female trips and 0.34 thousand male trips. Moving the same park to each gender's best location raises these figures to 6.60 and 1.60 thousand. Best-found three- and six-park layouts generate approximately 8.63–8.66 thousand female trips. In matched bootstrap draws, the observed network exceeds the central park by 3.92 thousand female trips, with an interval of $[1.82, 5.90]$, but falls short of the best-located single park by 1.61 thousand, with

an interval of $[0.68, 5.26]$. For the observed-versus-central comparison, the female advantage exceeds the male advantage by 3.11 thousand trips, with an interval of $[0.89, 4.72]$.

A centrally located park therefore performs quite differently from a park placed to maximise predicted trips. The best-found location also shifts across bootstrap draws. These comparisons concern predicted trip gains at candidate neighbourhood centroids, not precise parcels suitable for construction.

Reallocating the 315 small units under the nine-unit cap generates 8.05 thousand additional female trips and 1.87 thousand male trips. For women, the capped-small advantage over the best single park is 1.45 thousand trips, with an interval of $[-0.45, 1.68]$. Its difference from three best-found parks is -0.58 thousand, with an interval of $[-1.70, -0.02]$. The point estimates favour capped small units over a single well-placed park, although the bootstrap interval includes zero. Three well-placed parks generate somewhat more trips at this calibration.

Removing the cap raises the gains to 8.83 and 2.02 thousand trips while concentrating units in only nine neighbourhoods for women and eight for men. At $\rho = 1$, co-located units are equivalent to a single facility with their combined area. The gain from lifting the cap comes from greater flexibility in the spatial allocation of land, not the division of a facility into multiple units at the same location. Population-proportional allocation instead generates 4.62 and 1.08 thousand trips, slightly below the observed network. Population-proportional siting generates fewer predicted trips than the observed network and substantially fewer than the best-found layouts.

Utilisation and gender composition. Across common layouts evaluated for both genders, total net additional trips range from 1.41 thousand per day for one central park to 10.61 thousand for six best-found parks (Table E12). The female share varies less: it is approximately 81 per cent for the observed network, 76 per cent for the central park and 79–82 per cent across best-found layouts. The shares refer to net additional trips within the modelled purpose, not all visitors or people observed on surrounding streets.

Shares also differ across the alternative purpose definitions. At the main calibration and observed locations, women account for 81 per cent of additional RPTRA-composite trips, 87 per cent of broad non-work trips and 99 per cent of extended-composite trips (Table E9). These purpose categories overlap and do not represent counterfactual changes in facility services. Within the main RPTRA composite, shifting location affects total predicted additional activity much more than its female share.

Returns to size. Against one central park, the observed network breaks even at $\rho^{*f} = 1.317$ for women, with an interval of $[1.257, 1.347]$, and $\rho^{*m} = 1.244$ for men, with an interval of $[1.194, 1.287]$, at $\alpha = 0.05$. Within the scanned range $[0.5, 2]$, values below these crossings favour the observed network and values above favour the central park. A female crossing exists in every bootstrap draw and a male crossing in 88 per cent. The remaining male draws set the access coefficient to zero, leaving no trip-making difference. The male interval and the gender difference of 0.073, with an interval of $[0.039, 0.095]$, are conditional on the relevant crossings existing.

The thresholds differ when the large park is well placed. Against the best single park, the observed network's thresholds are 0.925 for women and 0.911 for men. For the capped reallocated small network, the thresholds are 1.054 and 1.043 against one best-located park and 0.980 and 0.987 against three best-found parks. Small layouts perform better below the reported crossings and large comparators above them. These point thresholds reselect locations at each ρ and do not have bootstrap intervals.

Appendix E.11 reports the calculations.

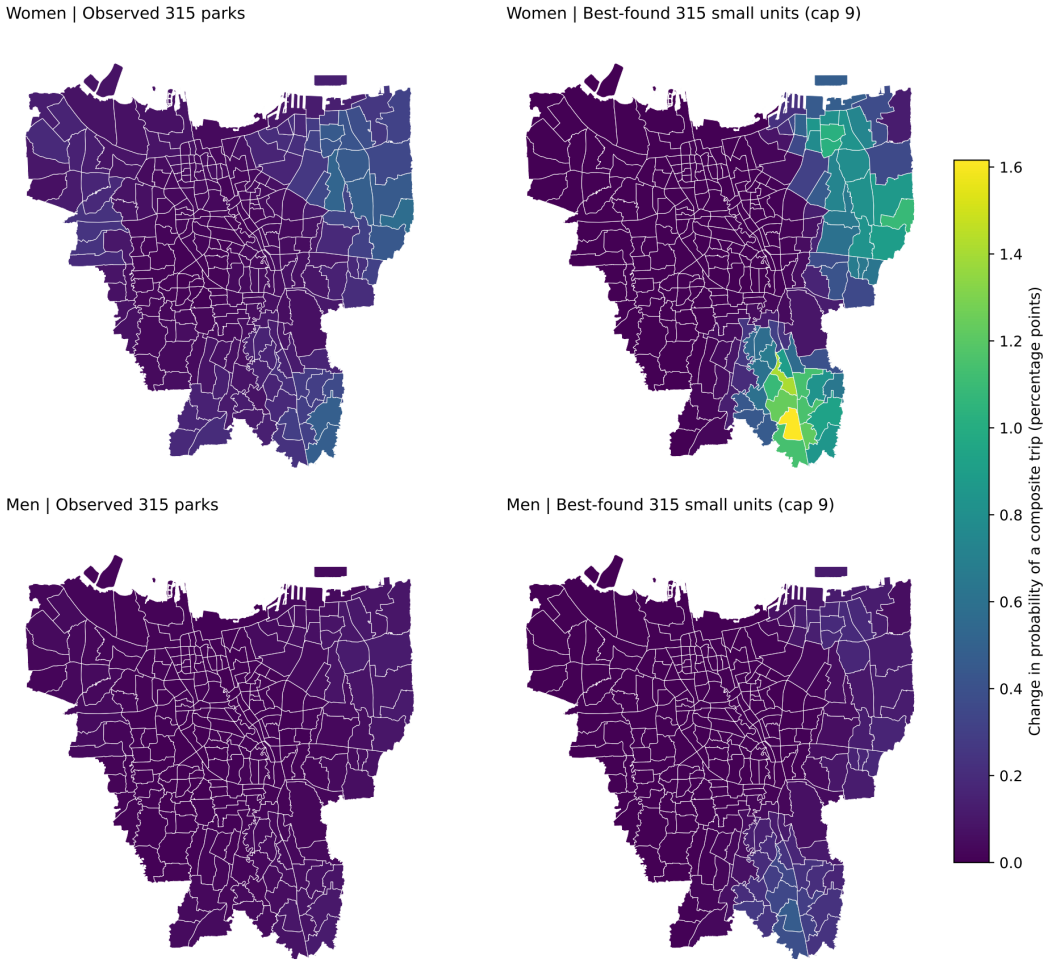
At $\alpha = 0.01$ and 0.10, the equal-attractiveness comparison still ranks best-located concentration above the observed network and central concentration below it. The observed-versus-central female thresholds are 1.290 and 1.338, and the male thresholds 1.226 and 1.260. Other rankings are less stable: at $\alpha = 0.01$, capped small units lose to the best single park even at $\rho = 1$. Alternative purpose definitions also yield different rankings and gender contrasts (Appendix E.5). The ranking of small and large facilities depends on the attractiveness calibration and assumed returns to area.

Where gains occur. Population and existing access jointly determine the geography of gains. A facility within reach of a dense origin with low access can generate substantial additional trip-making. Figures E2 and E3 compare the observed network with capped small-unit and three-park layouts. High-gain areas under the alternative layouts include peripheral neighbourhoods. The maps display changes in local trip-making probabilities, which differ from population-weighted citywide gains. These neighbourhood-level gains are conditional on candidate centroids and do not account for available parcels.

The larger predicted female trip gains and similar proportional street responses in Section 6 refer to different outcomes. A visit contributes to street occupancy only for the time its user is visible from the sampled road, and nearby counts include workers, passers-by and people waiting as well as facility visitors. The interior photographs do not supply a conversion from trips to street counts. I do not estimate a mapping from modelled trips to the visible-time outcome in the opening regressions. Better access may shorten travel, while longer travel or waiting can raise visible counts. Trip-making and visible street occupancy may consequently rank layouts differently.

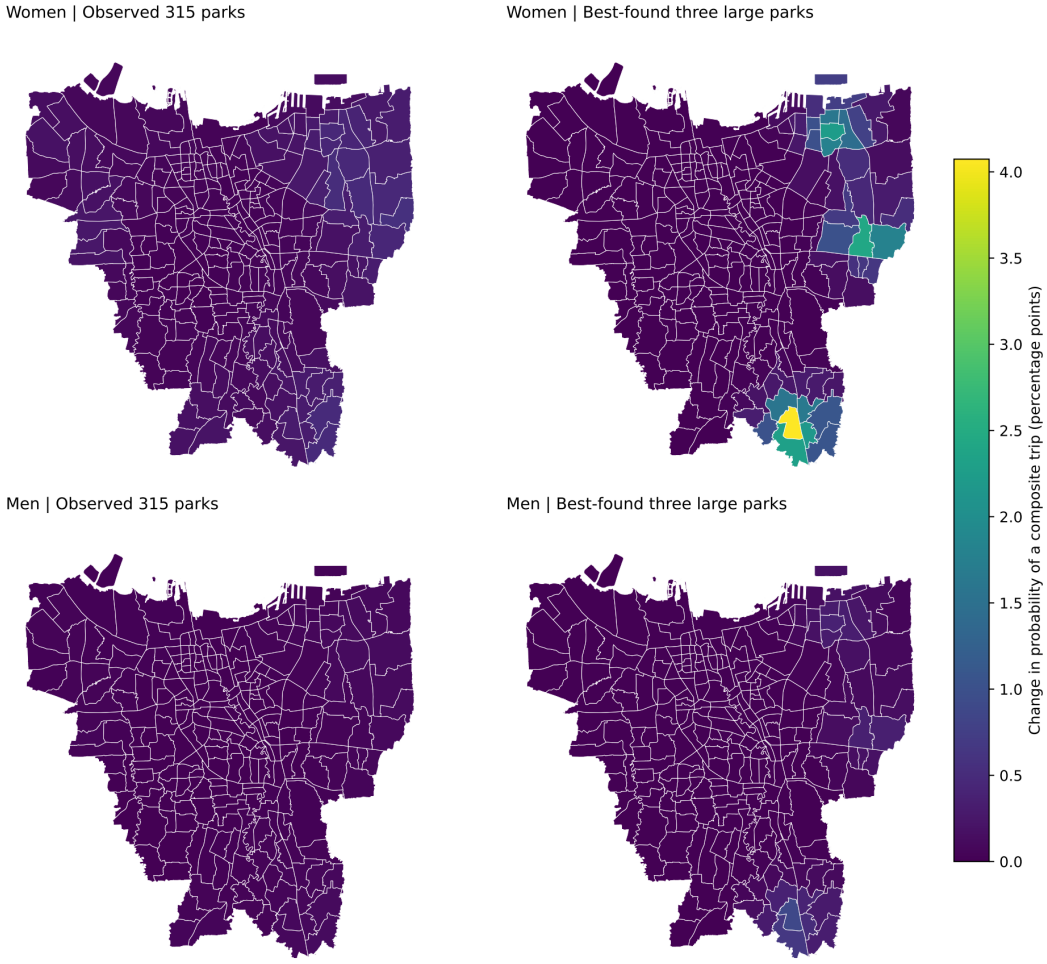
Equivalent distance. The expression $\ln[\Phi^g_{iq}(\mathcal{N})/\widehat{B}^g_{iq}]/\widehat{\delta}^g_q$ is the uniform reduction in effective distance to existing destinations that gives the same access increase. It is not actual travel saved. Since the denominator varies by gender, rankings in equivalent metres need not match rankings in additional trips.

Figure E2: Where trip-making increases: observed and reallocated small facilities



Notes: Model predictions combine 2002 SITRAMP travel estimates, fixed 2010 adult populations and the 315-site RPTRA register. The purpose combines recreation/sport/park trips and private trips to educational or medical destinations. All networks use 55.78 hectares and 261 candidate neighbourhood centroids. Gains are relative to the 2002 destination set with no added network. The calibration is $\alpha = .05$ and $\rho = 1$, so equal area gives equal total attractiveness. Facility attractiveness at the median site area equals α times median positive estimated destination attractiveness; ρ is its elasticity with respect to site area. Colours show changes in the probability of at least one qualifying trip on the survey day, in percentage points, by origin and gender. Left panels retain observed sites; right panels use best-found 315 equal-unit layouts with at most nine units per neighbourhood, selected separately to maximise each gender's net additional trips. Multi-site searches use greedy addition and single swaps, not proven global optima. All panels share a colour scale. These are point predictions without uncertainty shading; local probability changes differ from population-weighted citywide trip gains.

Figure E3: Where trip-making increases: observed sites and three large parks



Notes: Model predictions combine 2002 SITRAMP travel estimates, fixed 2010 adult populations and the 315-site RPTRA register. The purpose combines recreation/sport/park trips and private trips to educational or medical destinations. All networks use 55.78 hectares and 261 candidate neighbourhood centroids. Gains are relative to the 2002 destination set with no added network. The calibration is $\alpha = .05$ and $\rho = 1$, so equal area gives equal total attractiveness. Facility attractiveness at the median site area equals α times median positive estimated destination attractiveness; ρ is its elasticity with respect to site area. Colours show changes in the probability of at least one qualifying trip on the survey day, in percentage points, by origin and gender. Left panels retain observed sites; right panels use three equal-area large parks at their best-found locations, selected separately to maximise each gender's net additional trips. Multi-site searches use greedy addition and single swaps, not proven global optima. All panels share a colour scale. These are point predictions without uncertainty shading; local probability changes differ from population-weighted citywide trip gains.