Half-(head)way there: Comparing two methods to account for public transport waiting time in accessibility indicators

Half-(head)way there: Comparing two methods to account for public transport waiting time in accessibility indicators

Anson F. Stewart[1] and Andrew M. Byrd[2]

The Version of Record of this manuscript has been published and is available in Environment and Planning B 2022: https://doi.org/10.1177/23998083221137077

Abstract

Various methods have been developed to account for travel time variability and uncertainty when analyzing public transport networks and computing related accessibility indicators. In this paper we establish some convergence characteristics of one such method, implemented in the R5 routing engine, yielding guidelines for the minimum number of randomized schedules. This parameter has implications for result stability, analysis turnaround time, and computation costs. We also confirm that for travel time and accessibility results, there are spatially varying differences between our method and the conventional method relying on the assumption of half-headway waiting times. The conventional method appears to understate the benefits of transit in certain locations, particularly those served by multiple lines. Researchers and planning practitioners may find the R5 method preferable when analyzing complex networks or comparing transit scenarios where routes are specified in terms of headways or frequencies, rather than complete schedules with exact departure times for each trip.

[1]Conveyal LLC and MIT Department of Urban Studies and Planning
[2]Conveyal LLC

1 Introduction

Several recent articles describe new interactive sketch planning tools based on cumulative opportunities accessibility indicators (Conway et al. [2017], Conway and Stewart [2019], Stewart and Byrd [2021]). The open-source network analysis software underlying these tools, R5, has been widely reused (e.g. Pereira et al. [2021]) and favorably compared to other routing software in terms of features and performance (Higgins et al. [2022]). A key goal driving the development of R5 has been to characterize the variability and uncertainty in public transit travel times (and therefore accessibility) for riders departing at different times of day, where some or all transit routes are described in terms of headways or frequencies, rather than completely specified timetables. Such headway-based routes are a common way to represent long-term visions for transit network growth and change in scenario- and sketch-planning exercises.

The core methods in R5 treat travel times as random variables rather than single values. The underlying distributions can take on complex forms where multiple rides are chained together, or where multiple alternatives provide service to the same destination, particularly where headway-based sketch planning routes are overlaid onto a baseline network composed of fully timetabled routes. To address this problem, Conway et al. [2017] developed a Monte Carlo approach, sampling a large number of possible system-wide schedules meeting the scenario's headway constraints, while also examining all departure times in a multi-hour window at one-minute resolution.

This approach allows rapid computation of travel time distributions to millions of destinations from one origin, and when carefully optimized can provide region-wide accessibility figures from millions of origins in a matter of minutes. However, as a randomized sampling approach it does not exhaustively consider every possible combination of route schedules, and the results are expected to show variation between runs. The question then arises of how many randomized schedules, or Monte Carlo draws, must be examined to mitigate statistical noise and obtain stable, reliable results for use in a planning or decision process.

Conway et al. [2018] focus on quantifying the uncertainty in results and providing confidence intervals for this method. In everyday practice, however, constraints are placed on computation budget and turn-around time. It is preferable to allocate scarce computation resources to a higher number of draws yielding results with minimal error, rather than exactly characterizing the larger amount of error present with a lower number of draws. To make an informed trade-off between computation requirements and accuracy, we need practical guidelines for setting a sufficient number of draws before an analysis is begun.

In this paper we also compare two methods to account for waiting time on transit routes that lack completely specified schedules (i.e. routes without exact departure times specified for each trip):

The half-headway (hereinafter HH) method assumes a rider always waits half the headway of a route for a vehicle to arrive on that route. This conventional approach attempts to find a central tendency in total travel time by using the central tendency of wait time at each boarding event. For routes running at constant frequencies, waiting time does not vary at different departure times. This method is intuitive and perfectly appropriate for journeys that involve a single route. But it is decreasingly appropriate for journeys that could be made with multiple alternative routes, or that require transfers between routes.

The Monte Carlo (hereinafter MC) method in R5 evaluates many alternative system-wide schedules (draws), each of which coherently respects the headway parameters of every route in a scenario. Each MC draw is a combination of full timetables for the entire network. These timetables are provided by input data for schedule-based routes and generated randomly for headway-based routes. Waiting time and travel time are then derived for each draw and each minute of a specified departure window. When selecting a representative travel time, lower percentiles of the derived travel times represent both the rider choosing a better departure time from within a given window, and better timing of the network as a whole with respect to that rider's origin point and the spatial distribution of destinations.

MC is more complicated than HH, both to implement and interpret. Interpretability is an important criterion for accessibility indicators (Geurs and Van Wee [2004]); tradeoffs between ease of interpretation and quality of results should be weighed carefully. In this article we evaluate stability characteristics of the MC method and demonstrate that it provides better results than the HH method for journeys involving multiple routes, justifying its relative complexity.

This paper is structured around two topics: the convergence of travel time and accessibility values produced by the MC method as the number of randomized schedule draws increases; and second, the spatially varying differences in the indicator values produced by the MC method relative to the more conventional HH method.

1.1 Convergence

We expect travel time and cumulative-opportunity accessibility values to converge as the number of MC draws increases. More specifically, we expect that over a given number of trial repetitions with MC draws, the range and root-mean-square deviation (RMSD) of travel time for any specific origin-destination pair, and of accessibility at any specific origin, will be smaller with larger numbers of draws. We investigate whether results stabilize after a certain number of draws. If results indeed stabilize for a network known to present a high degree of uncertainty (e.g. arising from many frequency-based routes), the number of draws at which they stabilize could serve as a guideline for ensuring convergence in other, less uncertain networks. Finally, for a given number of draws we expect lower variation across trials in locations served by high-frequency routes that do not require transfers to access the destinations of interest.

1.2 Monte Carlo versus half-headway method

Compared to the HH method, we show the MC method yields shorter travel times and higher accessibility values for many origins in a real-world network, even when selecting travel times at or above the 50th percentile. But in peripheral areas where journeys involve multiple transfers and lower-frequency routes, the MC method yields longer travel times, suggesting that the HH method may misestimate travel times in a spatially biased way.

We expect the differences between these methods to be smaller for journeys served by a single high-frequency route, and larger for journeys served by multiple route alternatives with lower-frequency segments. This result relates to the common lines problem (Chriqui and Robillard [1975]), as discussed below.

1.3 Empirical case: Santiago de Chile

We test these methods using Santiago de Chile as a case study. Santiago's integrated transport network includes approximately 11,300 bus stops and 120 rail stations, served by the privately operated bus system formalized in 2007 (see Muñoz et al. [2014] for an overview) and Metro. Bus routes are loosely categorized as trunk routes, which traverse the city and often overlap on key corridors, and feeder routes, which are shorter and generally run with lower frequencies. Two characteristics of Santiago's land use and public transport make it a suitable for highlighting differences between the MC and HH methods.

First, jobs are highly concentrated, which makes cumulative opportunity measures of access to jobs especially sensitive to the selected travel time limit and other parameters. Garreton [2017] describes Santiago's evolution and spatial development: residential locations reflect a clear economic gradient; high-income households are clustered northeast of the city's historic core, and lower-income households tend to live in the southern and western sectors. Government programs in the 1980s resettled thousands of poor households from central areas to sprawling housing developments south of the city. Given the concentration of jobs in the historic core and along an axis northeast of Tobalaba (see Figure 1), today many low-income households endure long transit trips with multiple transfers to access jobs.

Second, Santiago's Metropolitan Directorate of Public Transport (DTPM) does not generally specify complete timetables; it instead publishes programmed route frequencies. The private operating companies responsible for dispatching buses have financial incentives based on route-level performance indices for frequency compliance and regularity compliance (Beltrán et al. [2013]). They therefore seek to operate individual route patterns with even headways, but there are no incentives for coordinated timetables between routes, even in corridors with overlapping routes. The lack of complete timetables implies a high degree of uncertainty across the network, which serves to clarify convergence characteristics, and differences between the MC and HH methods, of interest in this research.

In addition to these methodological reasons, transit unreliability is a salient issue in Chile. Protests against the transit system in 2019 spiraled into mass fare evasion campaigns, a state of emergency, and eventually a process for drafting a new national constitution. Underlying causes of this unrest include dissatisfaction with transit unreliability and decades of widening inequality (Rodríguez Mega [2019]). The analysis in our paper assumes all service operates as programmed, with deterministic headways and running times – holding fixed the foremost operational aspects of unreliability that could be captured by methods of estimating day-to-day variability in travel times and accessibility (e.g. Wessel and Farber [2019], Arbex and Cunha [2020]). Even assuming service operates as programmed, however, the uncertainty implied by the lack of complete schedules may play a role in perceived unreliability and users' dissatisfaction.

2 Past work

Typical transit-focused accessibility calculations and sketch planning applications rely on building a shortest-path tree from a selected origin to all destinations in a region. These trees are built by optimizing a single variable, such as travel time or generalized cost. Estimating waiting time in these methods is a challenge. Route choice models add nuance. But they are generally applied in assignment problems for a known origin and destination stop, rather than to build a tree from an origin to all destinations. Liu et al. [2010] offer a review of the extensive literature about such models.

2.1 Estimating waiting time

A "common approach" to estimating transit waiting times is the half-headway (HH) method (Curtis and Scheurer [2010]), used by Mamun et al. [2013], Farber and Grandez Marino [2017], and Swayne et al. [2018], among others. This approach considers the average wait time to be sufficiently representative of passenger experience. This may be the case for rides on a single headway-based route, but is decreasingly appropriate as alternative routes and transfers involving scheduled routes are taken into consideration.

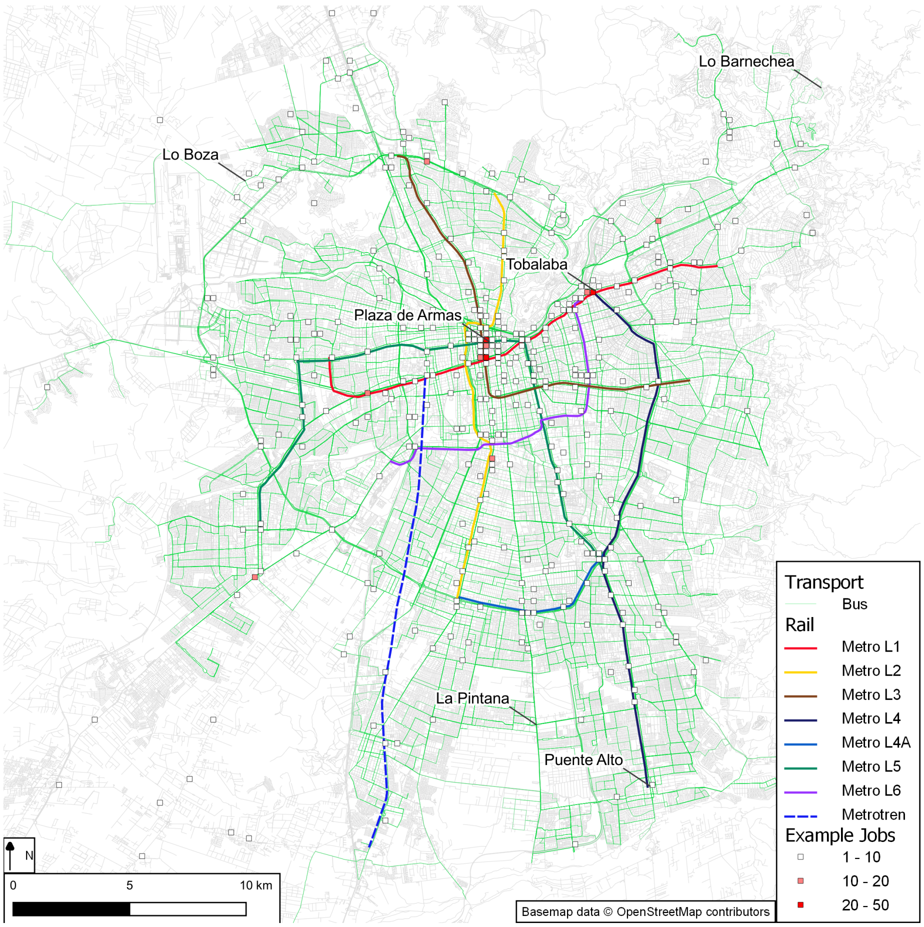

Figure 1: Public transport network and example job locations, Santiago de Chile.

When a transfer is required, the travel times reported to reach each destination are likely not all achievable within a single system-wide schedule; depending on the structure of the network, a schedule allowing riders to experience half-headway wait times at one point will often produce much worse wait times at some other points. The fact that a passenger boarded with half-headway wait, or minimal wait at one stop, may (together with the inter-stop travel times) necessarily mean that he or she cannot experience half-headway or better when transferring to another route at a downstream stop.

Finally, the half-headway assumption is problematic in the case of overlapping "common" lines in a corridor (see Chriqui and Robillard [1975]), or more generally in cases where multiple paths are available (as discussed in Conway et al. [2017]). Arriagada et al. [2019] evaluate rider behavior in the presence of common lines in Santiago, finding that passengers "use common lines as a unique alternative in 80% of the route sections" and that only 10% of passengers stick to one route when common lines are available.

To address the common lines problem, Arriagada et al. [2019] use hyperpaths. This approach identifies an attractive set of routes and calculates waiting time based on an assumed headway distribution for each route in the set, weighted by probability of boarding it. In the foundational work on hyperpaths, Spiess and Florian [1989] start with the half-headway assumption (p. 86), then extend to more general cases; they recognize that mechanisms for "constructing the links and the nodes of the generalized network depend very much on the particularities of the transit network and the degree of aggregation considered." In other words, a general method for enumerating attractive sets in large-scale networks may be computationally difficult. Li et al. [2015] estimate that solving an average all-to-one optimal hyperpath problem for the bus network in Chicago, which is similar in scale to Santiago (on the order of 10,000 stops), takes over 12 minutes of CPU time; they propose a heuristic alternative to identifying the attractive set, which reduces computation time substantially, but it is not guaranteed to be correct unless headways are assumed to follow an exponential distribution. In short, for rapid-turnaround, interactive sketch planning for large-scale transit networks, the assumptions and computation time that the hyperpath approach requires may be impractical.

Another strain of research in the transit assignment literature incorporates exact timetables – models for full schedule-based networks (the subject of Wilson and Nuzzolo [2008]), or mixed frequency- and schedule-based networks (Conway et al. [2017]). Such models can allow "adaptive path choice" Cats et al. [2011]. For example, Hickman and Wilson [1995] consider a "'clever' passenger" who "waits at the origin stop until a vehicle arrives and then makes a decision whether to board," using a case study with three schedule-based bus routes (albeit with stochastic headways and running times) and a frequency-based rail line.

2.2 Monte Carlo approach

Conway et al. [2018] also consider a "clever" passenger who must contend with variability due to departure time regardless of whether service is schedule-based or frequency-based, as well as uncertainty from under-specified service when routes are frequency-based. Such routes have a specified frequency but the offset of departure times from the top of the hour (*i.e.* phase) is treated as a random variable.

The MC approach handles these complexities, using a single parameter (the percentile of the overall travel time) to handle the variability due to departure time and the uncertainty due to phasing. Simulating possible system schedules, and sampling travel times and paths throughout a departure time window, obviates the need to enumerate attractive sets for hyperpaths *a priori*. Other advantages of the MC approach include: robust handling of services that only run during part of the departure time window, robustness against outliers, and ability to compare well-planned and ill-planned journeys and timetables. For a detailed example in a simple corridor, see the Supplemental Materials.

2.3 Measuring error due to sampling

Like our work here, Owen and Murphy [2019] also deal with the error characteristics of sampling methods in cumulative opportunities accessibility analysis over time windows. However, they focus on sampling *rider* departure times, while our work here focuses on sampling the space of all possible operational choices (route phases implying *vehicle* departure times) that satisfy a scenario with headway-based routes. We consider the problem of rider departure time sampling to be solved, as optimizations in R5 allow exhaustive examination. At one-minute resolution, there are only a few hundred distinct times in a morning peak period. In contrast, the space of all possible vehicle departure times grows exponentially with the number of routes and can easily reach the order of 10^{50}, necessitating a sparser sampling approach. Owen and Murphy [2019] use normalized root mean square error (NRMSE) to measure the quality of sampling strategies, applying a local indicator of spatial association (LISA) to identify discrepancies in proximity to transit infrastructure. Our measurement approach differs somewhat in its aggregation of multiple origin points and consideration of spatial patterns, but is generally comparable.

3 Methodology and data

Our overall approach follows Conway et al. [2017]: we build a routable baseline network from street data (from OpenStreetMap) and public transport schedules (from GTFS feeds), optionally modify or augment that network with scenario layers, calculate travel times to all destinations, select specific percentiles of those travel time distributions (often represented as isochrone curves), and derive location-based cumulative opportunities accessibility indicators at all origin points throughout the study region.

The walking, biking, or driving portion of the routing uses actual streets, and the public transit portion of the routing uses the actual published timetables when available. If published timetables are not available, as may be the case for routes with headway-based dispatching or for future scenarios at a sketch planning stage, we generate a large number of schedules satisfying all headway constraints in the MC approach, holding departure frequencies constant while randomly setting the phase of each route (i.e. the offset of its departure times from the "top of the hour").

This process is performed by version [GitHub branch/commit redacted to preserve author anonymity] of the R5 multi-modal routing engine, which requires certain parameters to be specified. For example, our analyses specify walking access/egress to/from transit stops, with a 20-minute limit per access/egress leg. We use a departure time window of one hour, calculating travel times for trips starting every minute within the window. We consider a range of travel time percentiles, including the 95th percentile – at this level, a commuter would complete a weekday morning commute late once every four weeks (see Furth and Muller [2006]).

3.1 Data and example locations for empirical case

Source data for the full case study include an OpenStreetMap extract [3] and the July 5 2019 version of GTFS published by Santiago's DTPM [4]. In this GTFS feed, all routes are frequency-based and lack explicit timetables except Metro Line 3 and the Metrotren, which for the purpose of this study are also modeled using headways rather than full timetables.

Population density data are from Chile's 2017 Census [5], The example jobs are derived from locations coded as the destinations of trips to work in Greater Santiago's 2012 origin-destination survey [6]. There are 966 example jobs, a relatively small sample that should heighten the sensitivity of cumulative opportunity job accessibility indicators to travel time fluctuations.

We selected six specific origins across the region for detailed results (Figure 1). The diverse network and land-use characteristics of these origins are described in the Supplementary Materials.

4 Results and discussion: travel time

This section discusses travel times from the six example origins to each cell in a 300m by 300m raster grid (30,687 destinations per origin). We conducted twenty trials for each origin, five percentile values, and five numbers of MC draws (60, 120, 480, 960, and 1200). These combinations yield approximately 92 million travel time values, or 5.5 billion when considering each departure minute separately. The statistics for convergence only include destination cells that were reachable in less than two hours in all trials. For comparisons with the HH method, we used the mean travel-time values across the 20 trials.

[3] See https://www.openstreetmap.org/

[4] Available online at https://openmobilitydata.org/p/dtpm-santiago-santiago/972/20190705

[5] See https://datosabiertos.ine.cl/dashboards/20568/censo-2017/

[6] See http://www.sectra.gob.cl/biblioteca/detalle1.asp?mfn=3253

4.1 Convergence

To illustrate travel time variability, isochrones were mapped from specific origins for 5th, 50th, and 95th percentile travel times. Large gaps between the 5th and 95th percentile isochrones indicate areas of high variability (i.e., where travel time can vary widely depending on specific departure time and vehicle operating schedules). As expected, the 5th and 95th percentile isochrones are farther apart for destinations connected to the origin by multiple infrequent alternatives requiring transfers. These maps also superimpose results from the 20 trials to illustrate the statistical noise arising from the MC method. As expected, the isochrones are much less “fuzzy” with 1200 trials than 60 trials, indicating a reduction in statistical noise. Maps and a detailed discussion of these results are available in the Supplementary Materials.

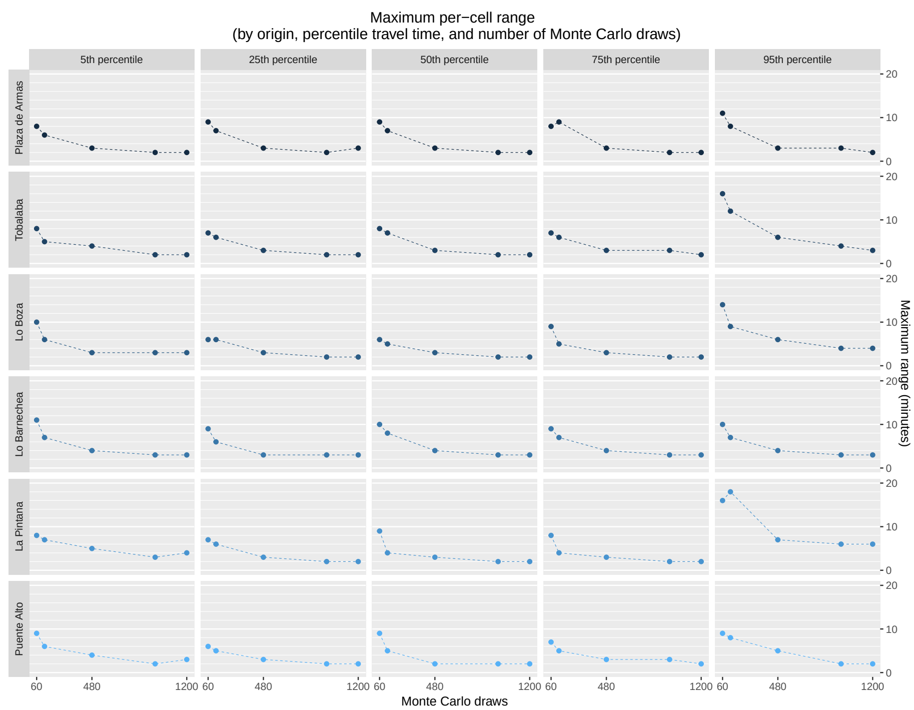

Figure 2: Maximum range of travel time

Figure 2 summarizes the maximum per-cell difference in travel time values for these six origins across 20 separate trials, each with the same number of draws. By 960 MC draws the travel time from any of these six origins to any destination does not differ by more than 3 minutes, except in a few cases at the 5th and 95th percentiles. Assuming deviations are symmetric, travel times at the worst-case destinations in these trials were within ± 1.5 minutes of the true limit value. In a few cases, the range does increase slightly when the number of MC draws is stepped up. This is not totally unexpected: even if the range is monotonically decreasing with a very large number of trials, any small finite number of trials is still varying around that limit, and that variance may be larger than the amount of convergence at the next higher number of MC draws.

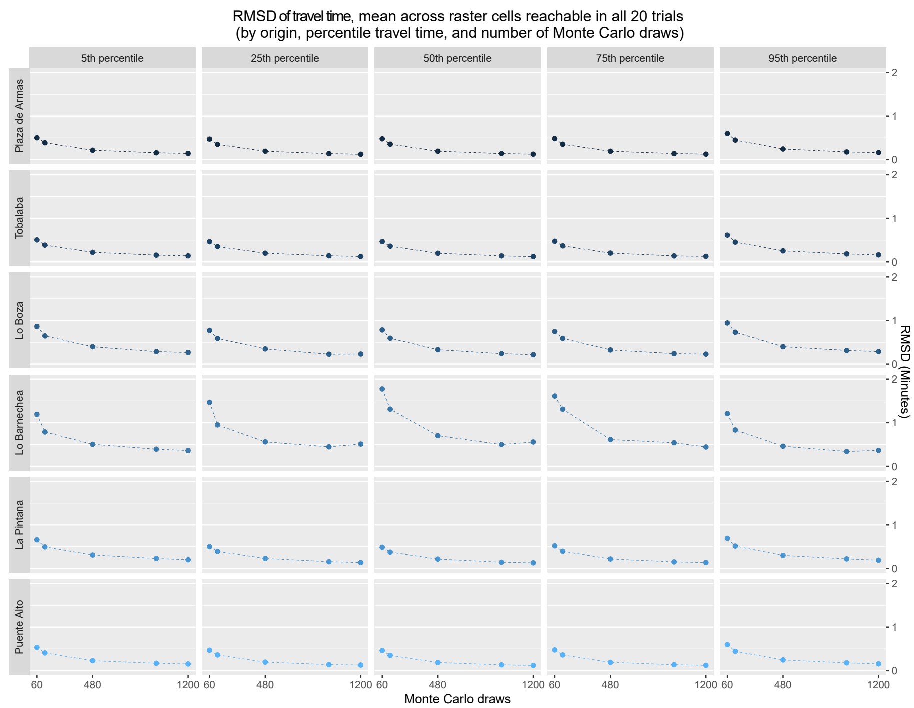

Figure 3: RMSD of travel time, mean across all reachable destination raster cells

Figure 3 summarizes the mean across reachable raster cells of root-mean-square deviation (RMSD) of travel time in the 20 trials. For all origins except one, RMSD is on average less than 30 seconds by 480 MC draws. The single low-frequency route (24-minute headways) serving Lo Barnechea makes it an outlier among the six example origins. RMSD for Lo Barnechea is generally higher than for other origins, and it peaks at 50th percentile travel time, while travel times from other origins tend to have higher RMSD at extreme percentiles.

4.2 Monte Carlo versus half-headway method

The previous results suggest reasonable convergence in travel time values by 1200 MC draws. In this sub-section, mean travel times for the 20 trials with 1200 MC draws are compared with travel times based on the HH method.

Figure 4 compares MC and HH travel times. The first row of this figure shows travel from an origin in the center of the region, Plaza de Armas. From this origin to most destinations, MC 50th percentile travel times are 1 to 5 minutes shorter than HH travel times. The differences between the two methods are less than a minute for destinations near stations along Metro Line 5; fast, frequent service from the Metro station adjacent to the origin means there are few other competitive routes from this origin, mitigating the common lines problem.

Results are mixed at 95th percentile travel times. At high travel time percentiles, the MC method reflects long waiting times, while waiting times for frequency-based routes in the HH method remain at half the headway. This difference explains why the upper right map of Figure 4 shows MC times approaching 10 minutes longer than HH times for travel to outlying areas served by single, low-frequency routes. On the other hand, some destinations for which there are many route alternatives (such as the Santa Rosa corridor running due south of downtown between Metro Lines 2 and 5), triggering the common lines problem, have 95th percentile MC travel times that are shorter than HH travel times.

Similar patterns are present for most of the remaining five origins. Again, the origin served by a single route with a 24-minute headway (Lo Barnechea) is an outlier. HH travel times reflect a 12 minute wait for this route; 5th percentile MC times reflect a very short wait (approaching 12 minutes shorter than HH), and 95th percentile travel times reflect a wait close to the full headway (approaching 12 minutes longer than HH).

Summarizing at a high level, most destinations shown in Figure 4 switch from blue (MC faster) to orange (HH faster) between the 75th and 95th percentiles, though there is substantial variability. This result is generally consistent with the simple corridor results discussed in the Supplementary Materials, where HH waiting times fell between the 75th and 95th percentile travel times.

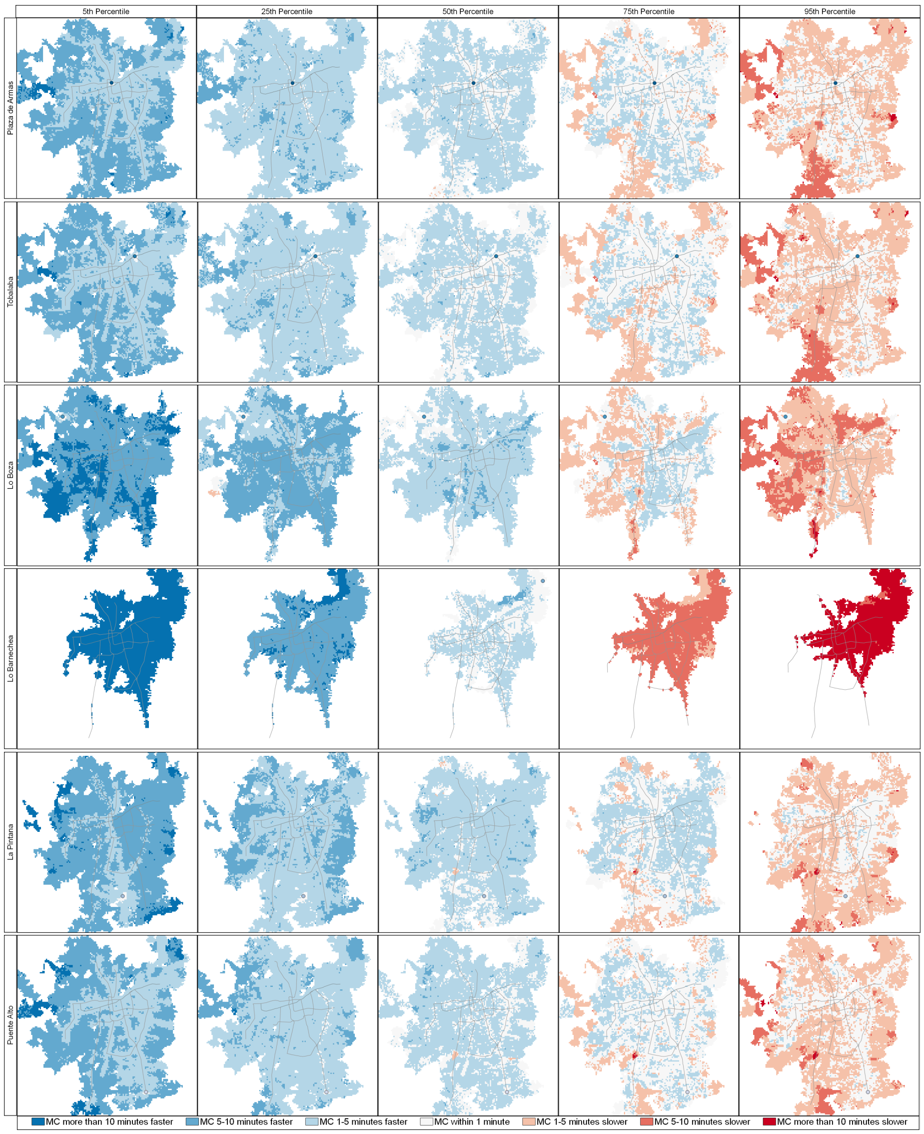

Figure 4: Difference in travel time, MC versus HH method, from six example origins at five travel time percentiles

5 Results and discussion: accessibility

Accessibility indicators are affected by the travel time considerations discussed in the previous section, as well as the interaction between the spatial distribution of opportunities and the chosen travel time limit or decay function. We expect accessibility values to converge, resulting in smaller ranges across multiple trials, as the number of MC draws increases.

As an initial indicator, we consider the number of example jobs reachable within 60 minutes at 50th percentile travel times. The spatial distribution of these 966 example jobs is shown in Figure 1. Maps in this section depict a number of jobs reachable from each origin in the region, using the 300m by 300m raster grid described above as origins.

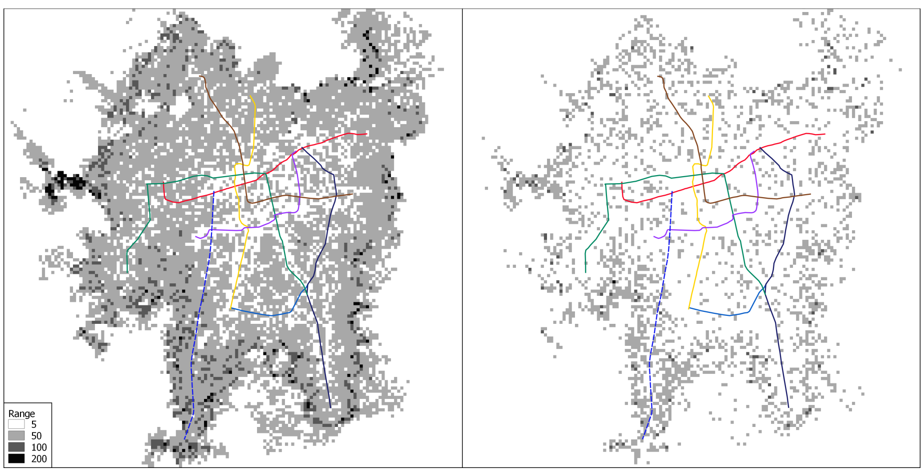

Figure 5: Range of accessibility values over 10 trials, for 60 MC draws (left) and 1200 MC draws (right)

5.1 Convergence

Figure 5 shows the maximum per-cell range in this accessibility indicator across 10 trials, using 60 MC draws (left) and 1200 MC draws (right). With 60 MC draws, the accessibility range for most cells is between 5 and 50 example jobs (0.5 and 5.0% of the total), and some cells have a range of approximately 200 example jobs (20% of the total). Assuming symmetrical distributions, this suggests that

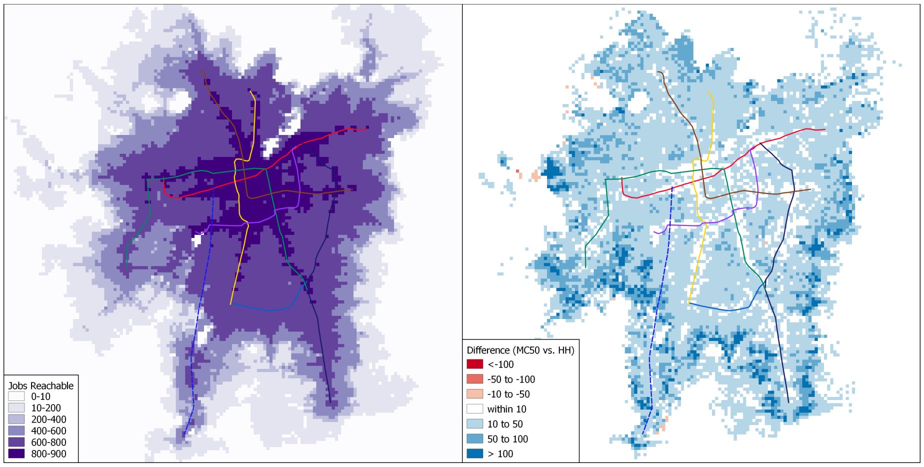

Figure 6: Example jobs reachable within 60 minutes given median travel times (mean of 10 trials, left) and difference from HH values (right)

results are with ± 10% of a stable value. With 1200 MC draws, the range for most cells is less than 5 example jobs (0.5% of the total).

In both cases, the origins with the largest range are close to the boundary of a 60-minute commute from the cluster of jobs downtown; the darkest cells in Figure 5 are generally coincident with the 60-minute contours from Plaza de Armas in Figure S1.

5.2 Monte Carlo versus half-headway method

The travel time results in Figure 4 suggest that for many trips in Santiago, MC median travel times are shorter than HH travel times. Given this result, the result in Figure 6 follows naturally: accessibility given median travel time is generally higher using the MC method. Again, the difference between the two methods is lower around Metro lines (corridors served by a single attractive route, mitigating the common lines problem) and higher at origins from which multiple lower-frequency route alternatives provide access to job centers.

6 Conclusion

Travel times converge on stable values within a computationally tractable number of draws. It follows that the accessibility indicator values derived from these travel times also converge, though accessibility converges more slowly at origin locations situated precisely at the accessibility indicator's travel time threshold relative to large clusters of opportunities. In this respect, the fringe of more slowly converging areas in Figure 5 shares conceptual similarities with the "ring of unreliability" in Cui and Levinson [2016]. We do not account directly for unreliability arising from poor schedule or headway adherence in actual operations, but the MC method can give some insight into variability that passengers may experience. Overall, the convergence characteristics described in this article confirm that the MC method is a practical approach for obtaining stable measurements of scenario impact without incurring excessive computing costs.

Travel times resulting from our MC method differ noticeably from those produced by the classic HH method, in a spatially varying fashion that a simple calibration factor would not capture. Accessibility indicators derived from the MC travel times also show perceptible differences relative to indicators derived from HH travel times, especially in areas served by multiple lower-frequency lines. This confirms the relevance of the more computationally demanding MC approach in estimating the impact of changes to complex networks where some or all lines are modeled in terms of frequencies rather than complete schedules.

6.1 Implications for practice

This research validates the original motivation for developing the MC approach: the classic HH approach understates the benefits of transit for certain places and overstates the benefits for others. Planners should question the validity of the half-headway assumption when making nuanced decisions between alternative scenarios specified in terms of headways, particularly where multiple lines or combinations of lines provide connections to areas of high opportunity density. In such situations, the HH method tends to overestimate waiting times, which artificially reduces the estimated accessibility benefits of transit. As common trunks and grid-like networks are widespread in both existing transit networks and prospective network redesigns, planners may find the MC approach tested in this research more suitable. Though more computationally intensive, a well-optimized implementation of MC still allows for near-instantaneous computation of travel times from one origin to all destinations, and can compute accessibility indicator values for every location in a metropolitan region in a matter of minutes.

This work provides initial empirically-derived guidelines for setting the number of draws or "simulated schedules," a user-specified parameter with implications for analysis turnaround time and computation costs. Satisfactory convergence of both median travel times and the derived accessibility indicators were achieved with 1200 draws; accessibility values converged to within ± 2.5% of the total number of example jobs, for almost all origins (see Figure 5). This result was obtained using a sparse, concentrated example set of jobs and a network with a high degree of uncertainty due to frequency-based specification of routes – both factors should hinder convergence. We therefore expect 1200 draws to be sufficient in other networks with generally similar frequencies, but less uncertainty (i.e. more routes represented with scheduled service).

This paper further characterizes the benefits of the approach described in Conway et al. [2017] for assessing accessibility in mixed schedule- and frequency-based networks. Santiago's DTPM has recently implemented a number of exact timetables for relatively low-frequency routes, including 16 that operate during the day and 21 that operate at night. The approach we describe could be useful in network planning as more of these routes are integrated with the frequency-based service in the rest of the network, or in cases where future scenarios are being developed.

6.2 Future research

Our results and conclusions suggest several avenues for further exploration.

Although we have determined a sufficient number of MC draws to achieve convergence on this particular transit network, in order to establish more general convergence guidelines more comparisons are needed between networks with different characteristics. We plan to perform a similar convergence analysis on networks with different ratios of frequency and fully scheduled lines, with higher and lower frequencies and differing network structures.

In place of the simple random sampling strategy used here, we intend to experiment with low-discrepancy sequences (quasi-Monte Carlo methods), which are expected to provide better error characteristics and more rapid convergence.

As mentioned in the conclusions above, convergence of accessibility indicators is slower where large clusters of opportunities are situated within the threshold isochrone's band of uncertainty. This problem is tied to the use of a hard-edged travel time threshold for inclusion of opportunities. Small shifts in travel time of even a few seconds can cause large clusters of opportunities (like large office buildings) to fall in and out of the isochrone from one trial to the next. This effect could be alleviated by applying a softer rolloff (e.g. gravity decay or sigmoid), effectively diffusing this localized error in the accessibility value across adjacent cells, over a distance controlled by the steepness of the threshold curve. We are currently integrating these changes and plan to compare RMS error figures for hard versus soft thresholds.

As Liu et al. [2010] paraphrase Hall (1982), "a transportation engineer can be very successful at reducing the travel time for ideal travellers, yet fail at improving the actual travel time seen by real travellers." Our work could be extended by integrating recent interdisciplinary research into characterizations of rider behavior and boarding strategies (e.g. Viggiano et al. [2014], Nassir et al. [2017], Ingvardson et al. [2018]).

Finally, the travel time distributions from which we select our chosen percentiles contain variation due to two separate factors: the departure time (which we vary uniformly across a multi-hour time window) and the uncertainty in vehicle arrival times. The latter is randomized in the MC approach but not the HH approach, while the former source of variation is present (and identical) in both the MC and HH approaches. It would be interesting to repeat these analyses with the departure time held constant, to completely isolate the effect of transit schedule randomization. However, arguably the true impact of different possible schedules (with their different implications for wait and transfer times) is not felt unless many different rider departure times are also considered.

References

R. Arbex and C. B. Cunha. Estimating the influence of crowding and travel time variability on accessibility to jobs in a large public transport network using smart card big data. *Journal of Transport Geography*, 85(January 2019):102671, 2020. ISSN 0966-6923. doi: 10.1016/j.jtrangeo.2020.102671. URL https://doi.org/10.1016/j.jtrangeo.2020.102671. Publisher: Elsevier.

J. Arriagada, M. Munizaga, A. Guevara, and C. Prato. Do passengers consider common lines? Observing and understanding route choice behaviour of public transport passengers from smart card data. In *TransitData 2019*, 2019.

P. Beltrán, A. Gschwender, and C. Palma. The impact of compliance measures on the operation of a bus system: The case of Transantiago. *Research in Transportation Economics*, 39(1):79–89, 2013. ISSN 0739-8859. doi: 10.1016/j.retrec.2012.05.026. URL http://dx.doi.org/10.1016/j.retrec.2012.05.026. Publisher: Elsevier Ltd.

O. Cats, H. N. Koutsopoulos, W. Burghout, and T. Toledo. Effect of Real-Time Transit Information on Dynamic Path Choice of Passengers. *Transportation Research Record: Journal of the Transportation Research Board*, 2217:46–54, 2011. doi: 10.3141/2217-06.

C. Chriqui and P. Robillard. Common Bus Lines. *Transportation Science*, 9(2):115–121, 1975.

M. W. Conway and A. F. Stewart. Getting Charlie off the MTA: a multiobjective optimization method to account for cost constraints in public transit accessibility metrics accessibility metrics. *International Journal of Geographical Information Science*, pages 1–29, 2019. ISSN 1365-8816. doi: 10.1080/13658816.2019.1605075. URL https://doi.org/10.1080/13658816.2019.1605075. Publisher: Taylor & Francis.

M. W. Conway, A. Byrd, and M. Van Der Linden. Evidence-based transit and land use sketch planning using interactive accessibility methods on combined schedule and headway-based networks. *Transportation Research Record*, 2653:45–53, 2017. ISSN 03611981. doi: 10.3141/2653-06.

M. W. Conway, A. Byrd, and M. Van Eggermond. Accounting for uncertainty and variation in accessibility metrics for public transport sketch planning. *Journal of Transport and Land Use*, 11(1): 541–558, 2018. ISSN 19387849. doi: 10.5198/jtlu.2018.1074.

M. Cui and D. Levinson. Accessibility and the Ring of Unreliability. *Transportmetrica A*, 2016. doi: http://dx.doi.org/10.1080/23249935.2016.1258441.

C. Curtis and J. Scheurer. Planning for sustainable accessibility : Developing tools to aid discussion and decision-making. *Progress in Planning*, 74:53–106, 2010. doi: 10.1016/j.progress.2010.05.001.

S. Farber and M. Grandez Marino. Transit accessibility , land development and socioeconomic priority: A typology of planned station catchment areas in the Greater Toronto and Hamilton Area. *Journal of Transport and Land Use*, 10(1):879–902, 2017.

P. G. Furth and T. H. J. Muller. Service Reliability and Hidden Waiting Time: Insights from Automatic Vehicle Location Data. *Transportation Research Record: Journal of the Transportation Research Board*, 1955(1):79–87, 2006.

M. Garreton. City profile: Actually existing neoliberalism in Greater Santiago. *Cities*, 65:32–50, 2017. ISSN 0264-2751. doi: 10.1016/j.cities.2017.02.005. URL http://dx.doi.org/10.1016/j.cities.2017.02.005. Publisher: Elsevier Ltd.

K. T. Geurs and B. Van Wee. Accessibility evaluation of land-use and transport strategies: review and research directions. *Journal of Transport Geography*, 12(2):127–140, June 2004. doi: 10.1016/j.jtrangeo.2003.10.005.

M. D. Hickman and N. H. M. Wilson. Passenger travel time and path choice implications of real-time transit information. *Transportation Research Part C*, 3(4):211–226, 1995. ISSN 0968090X. doi: 10.1016/0968-090X(95)00007-6. ISBN: 0968-090X.

C. Higgins, M. Palm, A. DeJohn, L. Xi, J. Vaughan, S. Farber, M. Widener, and E. Miller. Calculating place-based transit accessibility: Methods, tools and algorithmic dependence. *Journal of Transport and Land Use*, 15(1), Feb. 2022. ISSN 1938-7849. doi: 10.5198/jtlu.2022.2012. URL https://www.jtlu.org/index.php/jtlu/article/view/2012.

J. B. Ingvardson, O. A. Nielsen, S. Raveau, and B. F. Nielsen. Passenger arrival and waiting time distributions dependent on train service frequency and station characteristics: A smart card data analysis. *Transportation Research Part C: Emerging Technologies*, 90(March):292–306, 2018. ISSN 0968090X. doi: 10.1016/j.trc.2018.03.006. URL https://doi.org/10.1016/j.trc.2018.03.006. Publisher: Elsevier.

Q. Li, P. W. Chen, and Y. M. Nie. Finding optimal hyperpaths in large transit networks with realistic headway distributions. *European Journal of Operational Research*, 240(1):98–108, 2015. ISSN 0377-2217. doi: 10.1016/j.ejor.2014.06.046. URL http://dx.doi.org/10.1016/j.ejor.2014.06.046. Publisher: Elsevier B.V.

Y. Liu, J. Bunker, and L. Ferreira. Transit Users' Route-Choice Modelling in Transit Assignment: A Review. *Transport Reviews*, 30(6):753–769, 2010. doi: 10.1080/01441641003744261.

S. A. Mamun, N. E. Lownes, J. P. Osleeb, and K. Bertolaccini. A method to define public transit opportunity space. *Journal of Transport Geography*, 28:144–154, 2013. ISSN 09666923. doi: 10.1016/j.jtrangeo.2012.12.007. URL http://dx.doi.org/10.1016/j.jtrangeo.2012.12.007.

J. C. Muñoz, M. Batarce, and D. Hidalgo. Transantiago, five years after its launch. *Research in Transportation Economics*, 48:184–193, 2014. doi: 10.1016/j.retrec.2014.09.041.

N. Nassir, M. Hickman, and Z. Ma. Statistical inference of transit passenger boarding strategies from farecard data. *Transportation Research Record*, 2652(1):8–18, 2017. ISSN 21694052. doi: 10.3141/2652-02.

A. Owen and B. Murphy. Temporal sampling and service frequency harmonics in transit accessibility evaluation. *Journal of Transport and Land Use*, 12(1):893–913, 2019. doi: 10.5198/jtlu.2019.1379.

R. H. M. Pereira, M. Saraiva, D. Herszenhut, C. K. V. Braga, and M. W. Conway. r5r: Rapid Realistic Routing on Multimodal Transport Networks with R ^{5} in R. *Findings*, Mar. 2021. ISSN 2652-8800. doi: 10.32866/001c.21262. URL https://doi.org/10.32866/001c.21262.

E. Rodríguez Mega. What the protets and violence in Chile mean for Science. *Nature*, 575:265–266, 2019.

H. Spiess and M. Florian. Optimal Strategies: A New Assignment Model for Transit Networks. *Transportation Research Part B*, 23 (2), 1989.

A. F. Stewart and A. Byrd. Interactive Access for Integrated Planning (Chapter 17). In D. Levinson and A. Ermagun, editors, *Applications of Access*, pages 307–325. The University of Sydney, 2021. URL https://ses.library.usyd.edu.au/handle/2123/26890.

M. Swayne, M. Miller, M. Boarnet, and G. Painter. Innovation on Job Accessibility with General Transit Feed Specification (GTFS) Data: METRANS Project 16-08 Final Report. Technical Report December, METRANS, 2018. URL https://merritt.cdlib.org/d/ark%253A%252F13030%252Fm5s80dm9/1/producer%252FTier-1-UTC_16-08_Final-Report.pdf.

C. Viggiano, H. N. Koutsopoulos, and J. Attanucci. User Behavior in Multiroute Bus Corridors: Analysis by a Web-Based Survey. *Transportation Research Record: Journal of the Transportation Research Board*, 2418(1):92–99, 2014.

N. Wessel and S. Farber. On the Accuracy of Schedule-Based GTFS for Measuring Accessibility. *Journal of Transport and Land Use*, 12(1), 2019.

N. H. M. Wilson and A. Nuzzolo. *Schedule-based modeling of transportation networks*. Springer, 2008.

7 Supplemental Materials

7.1 Simple example illustrating the MC approach

A simple example helps illustrate the MC approach to the common line problem. Consider a trip in a corridor with three bus routes, each running every 30 minutes, with equal in-vehicle times between an origin and destination. If the relative phase of the routes yields evenly spaced departures, six effective departures (one bus every ten minutes) will be available to passengers in an hour; if routes are phased so that they depart simultaneously, passengers will have only two effective departures (three buses bunched together every thirty minutes).

In this simple example, waiting time is the only component of total travel time that varies with the specific departure time chosen, i.e. the exact time at which passengers depart their true origin to access transit. For a 60-minute long departure window, the 5th percentile travel time corresponds to the specific departure minute with rank 3, and the 50th percentile travel time corresponds to the specific departure minute with rank 30. If the timetables are perfectly out of phase (i.e. vehicles are evenly spaced), the 3rd shortest waiting time will be 0 minutes; passengers could in fact choose six departure times with 0-minute waits – one departure time that minimizes waiting time for each of the six effective departures. If the routes depart simultaneously (i.e. timetables are perfectly in phase), the 3rd shortest waiting time is 1 minute; passengers could wait 0 minutes for the first effective departure and 0 minutes for the second effective departure, but because there are no further effective departures in the one-hour window, the next best option is waiting 1 minute for one of the effective departures. Similarly, the rank 30 waiting time would be 5 minutes when out of phase, and 15 minutes when in phase.

More generally, the waiting time associated with these extreme cases of evenly spaced (out of phase) or completely bunched (in phase) can be calculated as:

$$\min(\lfloor \mathrm{rank}(w)/f_e \rfloor, \lfloor l/f_e \rfloor)$$

where $\mathrm{rank}(w)$ is the rank of the waiting time (e.g. 3rd best departure time), f_e is the number of effective departures corresponding to perfectly out-of-phase or in-phase timetables, and l is the length of the departure time window in minutes.

For networks exhibiting more complex combinations of frequency and phase across many routes, a number of effective departures may not be easy to calculate. Furthermore, if the phase parameters are random variables, the percentile of travel time parameter must reflect both variability with respect to departure time and uncertainty with respect to phase, which may not follow a well-behaved distribution. To capture the joint effects of variability and uncertainty, we generate many timetable combinations for the three routes, then measure the corresponding waiting times for passenger journeys starting at each minute over the 60-minute window.

The table below shows results for 6000 timetable combinations (Monte Carlo draws). As expected, the waiting times derived from the simulated schedules fall between the extreme out-of-phase and in-phase cases described above.

	Waiting time (minutes)		
Percentile	Out of phase	MC simulation result	In phase
5	0	0	1
25	2	3	7
50	5	6	15
75	7	11	22
95	9	19	28

Notably, the 50th percentile in-phase waiting time, which corresponds to the HH assumption (15 minutes), is greater than the median simulated waiting time (6 minutes). This result agrees with our expectation that the median waiting and total travel times from the MC method are shorter than the HH method when there are multiple path alternatives.

7.2 Description of example origins

7.2.1 Employment origins

The first three example origins represent employment locations. Travel times from these origins were calculated for a weekday departure time window of 7 PM to 8 PM, the last hour of the evening peak in Santiago. This window is appropriate for work-to-home commutes and, assuming symmetric AM and PM peak service, calculating indicators of employers' access to employees.

Plaza de Armas Historic city downtown and employment hub, with nearby service including frequent Metro (Lines 1, 2, 3, 5) and a dense grid of trunk bus routes radiating to other parts of the city.

Tobalaba Center of an axis running northeast from downtown that has seen "high-rise residential and business densification" [Garreton, 2017, Fig. 9], with nearby service including frequent Metro (Lines 1, 4, 6) and many trunk routes along the city's main northeast-southwest corridor

Lo Boza Outlying industrial park northwest of downtown, with nearby service provided by five circumferential bus lines

7.2.2 Residential origins

The remaining example origins represent residential locations. Travel times from these origins were calculated for a weekday departure time window of 7 AM to 8 AM, the first hour of the morning peak in Santiago. This window is appropriate for home-to-work commutes and calculating indicators of residents' access to jobs.

Lo Barnechea Residential neighborhood in the northeast of the city, with nearby service provided by a single feeder bus route, C13, which has a 24-minute headway during the morning peak.

La Pintana A center of "low-rise social housing [and] discontinuous densification" [Garreton, 2017, Fig. 9], to which many poor families were relocated by the military dictatorship in the 1980s, with nearby service provided by five primary north-south trunk routes

Puente Alto A "conurbated town" [Garreton, 2017, Fig. 9], with nearby service including Metro Line 4, three primary trunk routes, and many local feeder routes

7.3 Additional travel time results

Figure S1 shows 30- and 60-minute travel time contours from Plaza de Armas, at the 5th, 50th, and 95th percentiles of travel time to each destination. As seen in the top row of the figure, the vast majority of the region's residents are within a 60-minute commute from this central location. The distance between the contours for the 5th and 95th percentiles of travel time is, in most places, well under 1 km, reflecting the low variability in travel times given the plethora of frequent transit options serving this origin. Contours for all 20 trials are superimposed, as seen more clearly in the zoomed insets shown in the bottom row of the figure. With only 60 MC draws (1 simulated set of schedules for each minute in the departure time window), there is substantial noise in the results – note the "fuzziness" in the contours northwest and southeast of the Quilín station on Metro Line 4. With 1200 MC draws, this noise is substantially reduced, as the travel time values have converged.

Figure S2 shows similar travel time contours from Lo Boza. Along some corridors, the distance between the contours for the 5th and 95th percentile travel times exceeds 4 km, reflecting a much higher variability arising from the relatively low frequency of the circumferential bus lines. The 60-minute contours suggest various route alternatives for travel from Lo Boza toward downtown – south via bus then east via Metro Line 5, or east via bus then south via Metro Line 2 or 3. Using 95th percentile travel times, the relatively low frequencies of the bus routes imply long waiting times, making the Metro lines barely reachable within 60 minutes of total time. Using 5th percentile travel times, for "clever" passengers who can adjust their departure times to minimize waiting, there are options to reach downtown within 60 minutes of total time. Even given deterministic headways and running times, passengers who start their trips from this origin at a specific time every day may see substantial variability in how far they can travel depending on whether that time happens to be convenient (e.g. 5th percentile) or inconvenient (e.g. 95th percentile) relative to any given day's schedule offsets. Last, comparing the left and right of this figure shows a substantial reduction in "fuzziness" when using 1200 draws.

Figure S3 shows travel time contours from La Pintana, superimposed on a map of example job locations. This figure helps illustrate how the statistical noise from the MC method could affect cumulative opportunity accessibility indicators. If "fuzziness" in the contours overlaps locations with many jobs (e.g. the 30-minute, 5th percentile contour northwest of the origin, or the 60-minute contours

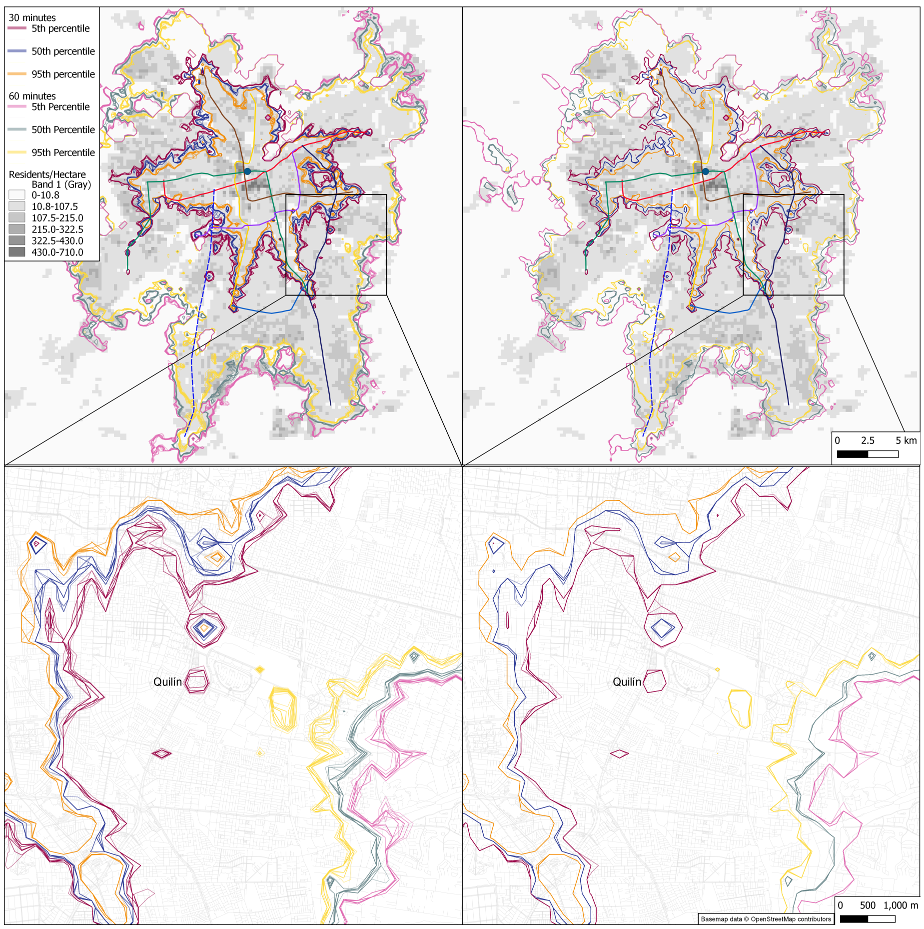

Figure S1: Travel time contours from Plaza de Armas, using 60 MC draws (left) and 1200 MC draws (right). Basemap © OpenStreetMap contributors

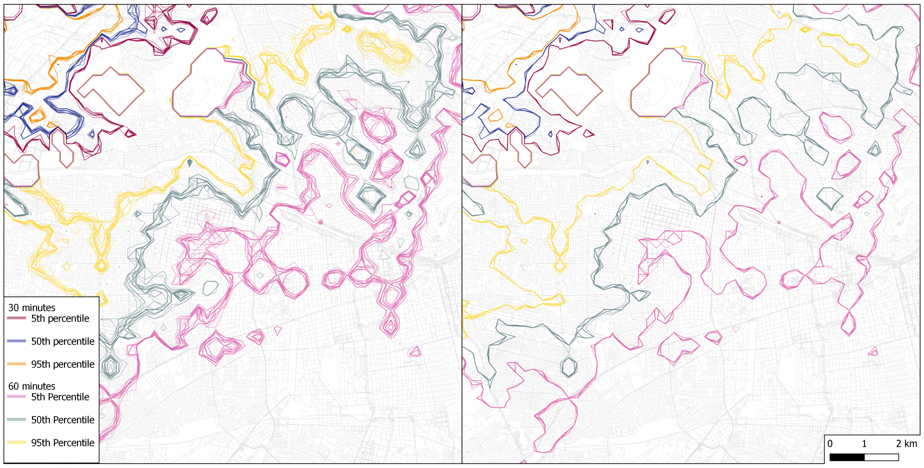

Figure S2: Travel time contours from Lo Boza, using 60 MC draws (left) and 1200 MC draws (right). Basemap © OpenStreetMap contributors

around Plaza de Armas, in Figure S3(left)) a cumulative opportunity job accessibility indicator would be especially unstable between trials.

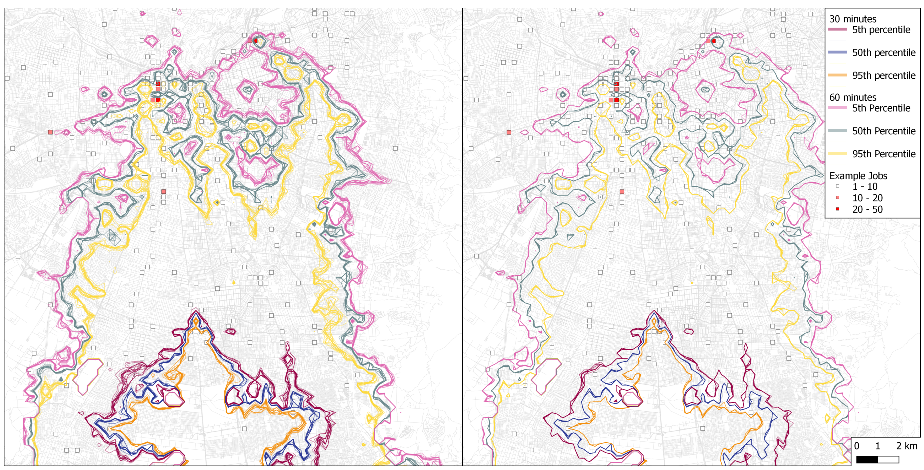

Figure S3: Travel time contours from La Pintana, using 60 MC draws (left) and 1200 MC draws (right), overlaid on example jobs. Basemap © OpenStreetMap contributors