Mobile Network Coverage and Road Surface Condition as Joint Constraints on Digitally Booked Transport Beyond State Capitals

Oladimeji Basit Alaka[1,*]; Oluwabusola Ajayi[2]; Olayemi Bolaji[3]
[1]Rivetline Concept Limited, Lagos, Nigeria
[2]Bisaj Engineering Limited, Osogbo, Nigeria
[3]Leedigital Research and Innovation, Ibadan, Nigeria
[*]Corresponding author: basitoladimeji@gmail.com
DOI: 10.56201/ijemt.v8.no5.2022.pg142.168

Abstract
A digitally booked ride needs two things at once: a signal to make the booking and a road the vehicle will travel. Telecommunications statistics measure the first, road statistics the second, and neither measures the pair. We build both layers for 958 settlements of 500 people or more beyond the capitals of Osun and Kwara, 5.18 million residents, from open cell records and 37,905 km of OpenStreetMap centerline. Condition is imputed for 53.3 percent of that length, and the layer's two independent readings agree at a Cohen's κ of 0.041, so no settlement's grade is reliable and no correlation computed on it is separable from that error. We expected the layers to correlate. They do not, at a Pearson correlation of 0.051 and a permutation $p = 0.487$, but a null this noisily measured is uninformative. What survives needs no correlation. The excess of the union over the worse single-sector statistic is identically the population that fails coverage and passes condition: 748,898 people, or 14.46 percentage points at the base thresholds. It is positive in 55 of the 56 cells of a radius-by-threshold sweep and zero in the 56th, between 0.00 and 18.85 points. A second sweep using distance to the nearest sealed classed road narrows that margin to 0.00 to 6.55 points, exactly zero in 9 of 42 cells, so the correction is real in direction and not portable in size. Half the population beyond the capitals fails at least one test, and the two failures fall on different people only under population weighting.

Keywords: *digital divide, road condition, joint constraints, ride-hailing, accessibility, OpenStreetMap, mobile coverage, Nigeria*

1 Introduction

Digitally booked transport arrived in Nigerian cities during the last decade and is now ordinary in Lagos and Abuja (Giddy 2019; Alcorn and Karner 2021; Olowosegun et al. 2021). Its edge is sharp and it is documented. By the end of 2021 Bolt served both of the state capitals in this study, listing Ilorin by December 2020 (Bolt Technology OÜ 2020) and Osogbo by October 2021 (Bolt Technology OÜ 2021a), while Uber operated in exactly five Nigerian cities, Lagos, Abuja, Benin City, Port Harcourt and Ibadan, and neither Osogbo nor Ilorin was among them (John 2021; Kuola 2021). Bolt was the more expansive of the two, listing 33 Nigerian cities in October 2021 (Bolt Technology OÜ 2021b), so the edge is not a matter of one operator being small. Within Osun and Kwara those 33 cities are the two capitals and nothing else. The service stops at the capitals in these two states, and what lies beyond them is the subject of this paper.

The usual explanations for that edge divide along sectoral lines. One is connectivity: Adeleke

(2021) and Odusanya and Adetutu (2020) document the Nigerian digital divide, and International Telecommunication Union (2021) and GSMA (2021) the regional picture. The other is roads: World Bank (2019c) puts rural access at 57.3 percent in Osun and 21.7 in Kwara, and Porter (2014) documents what poor surfaces do to service provision.

Both explanations are correct and neither is sufficient, because a booking requires both conditions simultaneously. A settlement with a strong signal and an impassable approach road cannot be served, and neither can a settlement with a good road and no signal. The excluded population is therefore the union of the two failures, and no published statistic measures a union, because the two layers are held by different ministries and reported separately (Federal Ministry of Communications and Digital Economy 2020, 2019; World Bank 2019a).

We began with a hypothesis about that union. If coverage and road condition are positively correlated, as one would expect if both follow economic development, then the two failures overlap heavily and the union is smaller than treating them as independent would suggest. If they are uncorrelated or negatively correlated, they compound and the union is larger. The size of the excluded population would then depend on a correlation nobody has measured.

The hypothesis is wrong, and Section 5 reports the null in full. The observed union differs from the independence prediction by an amount that is neither statistically nor practically distinguishable from zero and whose sign reverses under a modest change of threshold. Nothing is built on that null, because the condition layer carries an agreement statistic of $\kappa = 0.041$ and a correlation measured through that much noise is pulled toward zero whatever the underlying association is. Section 7 sets out why the null cannot be read as evidence about anything but the layer. What survives is a comparison against the margins rather than against independence, and it survives precisely because it never asks how the two layers covary. The union exceeds the worse of the two single-sector statistics by 14.46 percentage points at the base thresholds and by a positive amount in 55 of the 56 cells of the sweep. That is the operational point: whichever ministry has the worse number, the excluded population is larger than that number, because part of it is failing a test that ministry does not administer.

The setting supports that measurement, and every property of it that makes the measurement possible is also a constraint on what the paper can claim:

- ***Both layers can be built from published data.*** Cell site locations, road geometry and gridded population are all published, so the joint distribution can be built without any proprietary input (OpenCellID, n.d.; OpenStreetMap contributors, n.d.; WorldPop 2020).
- ***The thresholds are not arbitrary***. The 5 km service radius we adopt is the Universal Service Provision Fund's own published definition, so it is inherited rather than set by us.
- ***The capitals must be excluded***. Osogbo and Ilorin are served on both layers and would dominate any population-weighted statistic, which is why this study is explicitly about what lies beyond them.
- ***Road condition tagging is incomplete.*** Some combination of the tags covers 46.7 percent of network length and the remaining 53.3 percent has to be imputed, so the condition layer requires bounding rather than a point estimate (Barron et al. 2014; Barrington-Leigh and Millard-Ball 2017). What the paper contributes follows from that. We build and publish the joint distribution of mobile coverage and road condition for 958 settlements beyond two state capitals, entirely from open data. We test the correlation hypothesis, report the null with its sweep, and set out why the null is uninformative given the condition layer's measured unreliability. And we establish the result that does hold: the union of the two constraints exceeds either sector's own worst statistic by a margin large enough to change who counts as excluded, positive in 55 of 56

threshold pairs under the paper's condition test and in 33 of 42 under an alternative one, though its size ranges over an order of magnitude across both sweeps.

Section 2 reviews the connectivity, road condition and compound deprivation literatures. Section 3 describes the two layers, Section 4 the frame and the measures. Section 5 reports the marginals, the correlation null, the constraint classes and the union. Section 6 reports the sweeps and the imputation bounds, and Section 7 what follows.

2 Related Work

2.1 Measuring connectivity, and the gap between claimed and actual

Rural connectivity is measured badly and known to be. Federal Communications Commission (2019) documents the gap between operator-claimed coverage and measured coverage in a well-resourced regulatory environment, and Mangla et al. (2021) and Riddlesden and Singleton (2014) examine what open measurement datasets can and cannot substitute for it. For Nigeria, Adeleke (2021) and Odusanya and Adetutu (2020) document the divide, Okoro and Iwuji (2019) and Faruk et al. (2013) the propagation and quality picture, and Bello et al. (2016) looks at Kwara specifically. Salemink et al. (2017) reviews rural digital development, Houngbonon et al. (2021) and Bahia et al. (2020) the infrastructure and welfare links, and Rodríguez-Castelán et al. (2021) adoption itself. Oughton (2021) treats the policy design problem.

This literature measures one margin very carefully.

2.2 Road condition, and assessing it cheaply

Sayers et al. (1986) and Eaton et al. (1987) established the roughness and unsurfaced road condition frameworks still in use, and Mbabazi (2019) measures the operating cost consequence in an African setting. World Bank (2019c) and World Bank (2019a) give the Nigerian rural access picture, and Porter (2014) what poor roads do to services.

On cheap assessment, Griew et al. (2013) and Kelly et al. (2013) use imagery for street audits and both report agreement statistics for the surface-quality items specifically, which is the closest published analogue to what we compute here. The agreement machinery itself is Cohen (1960), whose κ we use unweighted and with linear and quadratic weights. The one directly on-point source, Bianchini et al. (2010) on inter-rater reliability of manual pavement distress evaluation, could not be obtained: every route returned HTTP 403, neither body nor abstract was read, and no value is quoted from it here. It is named because a reader checking this paper's agreement result should know where the carriageway benchmark would come from, not because it informed the method.

This literature measures the other margin, equally carefully and entirely separately.

2.3 Compound disadvantage

The multidimensional poverty tradition has thought hardest about what it means to fail on more than one dimension. Alkire and Foster (2011) and Alkire et al. (2015) set out the counting approach, Atkinson (2003) the multiple deprivation framework, and Ravallion (2011) the case against composite indices. In transport, Lucas et al. (2016), Church et al. (2000) and Kenyon et al. (2002) develop transport poverty and social exclusion, Velaga et al. (2012) the rural case, and Geurs and Wee (2004) the accessibility measurement frame.

The counting approach is the closest existing tool, and the difference here is that our two dimensions are not two aspects of deprivation to be weighted together but two independent preconditions for a single service, either of whose failure is disqualifying. That makes the union

the quantity of interest rather than a weighted index.

2.4 The data this study is built on

Haklay (2010), Barron et al. (2014), Barrington-Leigh and Millard-Ball (2017), Mahabir et al. (2017) and Herfort et al. (2021) assess OpenStreetMap completeness and bias, which bears directly on the 53.3 percent of our network carrying no condition tag. Anwar and Graham (2020) examines the hidden geography of platform work in Africa.

2.5 What is missing

No study we located builds both layers at settlement level for the same population and reports their joint distribution. The connectivity literature and the road literature each measure their own margin well, and the compound deprivation literature supplies a framework designed for weighting rather than for union. That combination is the gap.

3 Data

Table 1 lists every input with its source, license, vintage and resolution.

Table 1. Open data layers used to build the two constraint surfaces, with the vintage of the data content and the date it was retrieved recorded separately.

Layer	Source	Content vintage	Retrieved	Resolution	License
Cell sites	OpenCelliD, MCC 621	Oct 2016	14 Mar 2022	point, per cell	CC BY-SA 4.0
Mobile speed	Ookla Speedtest	2021 Q4	14 Mar 2022	*z*16, ≈480 m	CC BY-NC-SA 4.0
Road network	OpenStreetMap	live database	14–15 Mar 2022	vector, way level	ODbL 1.0
Population	WorldPop	2020 estimate	15 Mar 2022	3 arc-sec, ≈100 m	CC BY 4.0
Boundaries	GADM 4.1 level 2	2022 release	15 Mar 2022	vector, LGA	academic use
Place names	OSM place=*	live database	14 Mar 2022	point and area	ODbL 1.0

Note. MCC, mobile country code; LGA, local government area. Road geometry and place names are retrieved through the Overpass API; the WorldPop raster is the UN-adjusted constrained product; the Ookla tiles are web-mercator zoom 16. Sources are OpenCellID (n.d.; Ookla, n.d.; OpenStreetMap contributors, n.d.; WorldPop 2020; GADM, n.d.); the Ookla tiles are licensed CC BY-NC-SA 4.0 and are used here under that attribution. The OpenCelliD extract holds 30,126 records inside the study bounding box, reducing to 19,925 distinct operator-sites on a 250 m grid per operator, of which 8,743 carry a data-capable radio and 11,182 carry GSM only; inside the two state polygons themselves the counts are 5,785 operator-sites and 2,516 data-capable ones. The Ookla extract holds 913 tiles built from 2,694 tests on 1,185 devices. **The cell layer's content vintage is the column that matters and it is not 2021.** The value entered in that column is a median record date rather than a release date, because OpenCelliD is a continuously edited register and issues no version. No record post-dates the study period, but the median updated timestamp inside the box is 4 October 2016, only 0.35 percent of records were touched in 2019 or later, none in 2021, and the latest record of any kind is 14 August 2020. The road network is a snapshot of a continuously edited database and no earlier extract exists in the

cache.

3.1 The coverage layer

Cell site locations come from OpenCelliD (OpenCellID, n.d.), from which we took 204,384 Nigerian records. Within the study bounding box, which is drawn generously so that a mast just outside a state line still shortens a settlement's distance, 30,126 records reduce to 19,925 distinct operator sites after deduplication, of which 8,743 are data-capable rather than voice only. Those two counts are bounding-box counts and they spill into Oyo, Niger and Kogi. Inside the two state polygons the register holds 5,785 operator sites, 2,516 of them data-capable, on 3,618 distinct mast locations, and that is the figure to compare against any state-level count.

The data-capable layer is in practice a third-generation layer. Of the 30,126 in-box records, 17,937 are GSM, 12,177 are UMTS, twelve are LTE and none is NR. Twelve LTE records across two states cannot support any statement about fourth-generation reach, so where the national series reports 37 percent fourth-generation coverage below, that figure is context this study cannot test, confirm or contradict. Everything the paper calls data-capable is UMTS.

The layer's effective vintage is 2016, not 2021. The extract was filtered so that no record post-dates the study period, which protects against anachronism, but the binding fact runs the other way: the median record inside the box was last updated on 4 October 2016, 0.35 percent were touched in 2019 or later, and the most recent record of any kind dates from August 2020. A 2016 snapshot used to characterize 2021 service misses five years of mast building, so it understates coverage, overstates the coverage-failing population, and inflates both the union and the margin the paper reports. Nothing corrects it, because no other snapshot is cached.

A cell database is not a coverage map, and we treat it as a proxy with a stated radius. The 5 km service radius is the Universal Service Provision Fund's own published definition, which means the threshold is inherited rather than chosen. Because GSM propagates further than the data-capable technologies, we report a data-capable layer and an any-radio layer side by side rather than only the one the headline uses. Across the whole population of the two states, 79.3 percent live within 5 km of a site of any radio and 90.1 percent within 10 km, against 69.9 and 84.0 percent for the data-capable layer. Inside the 958-settlement study frame the same gap is wider, 11.3 points or 586,496 people who are reachable by telephone and not by application, and a further 805,085 people have no site of any radio within 5 km at all. The any-radio layer at 5 km therefore understates the published 89 percent second-generation figure by 9.7 points, which is what a single radius applied to two propagation regimes must do.

Ookla open measurements for the fourth quarter of 2021 (Ookla, n.d.), 913 tiles carrying 2,694 tests, are an independent presence check on the site-derived layer and never the layer itself, because their spatial coverage is far too sparse for a settlement-level map. The check returns 67.6 percent agreement on presence, computed across all 960 clusters including the two capitals rather than on the study frame: 608 show neither a data-capable site within 5 km nor a Speedtest tile, 41 show both, 301 show infrastructure with no test run, and ten show a test with no registered site. The 301-cluster cell is not a disagreement, because a tile exists only where somebody with a smartphone ran a test, so its absence is evidence about usage rather than about infrastructure. Only the ten-cluster cell is a genuine conflict, and at 1.0 percent it is small. Where tests exist the median download rate is 11.7 Mbit/s, well above what a booking application needs, which supports the framing that the binding constraint beyond the capitals is presence rather than speed.

3.2 The condition layer

Road geometry and condition come from OpenStreetMap (OpenStreetMap contributors, n.d.), an Osun extraction with six Kwara tiles added, giving 37,905 km of motorable centerline. Condition is read from the surface, tracktype and smoothness tags, of which some combination covers 46.7 percent of network length. Completeness is strongly ordered by class, from 96.4 percent of trunk length down to 17.3 percent of residential and 11.6 percent of track, so the untagged half is concentrated in exactly the local streets a pickup would use.

The remaining 53.3 percent is imputed, and we bound it in both directions throughout rather than reporting a point estimate. Section 6 gives the bounds.

3.3 The frame

Settlements are contiguous clusters of populated WorldPop cells holding at least 500 people (WorldPop 2020), giving 958 settlements and 5,179,339 residents once the two state capitals are removed. Table 2 reports the frame and Figure 1 maps it.

The frame's most consequential property is its size structure. Sixteen settlements of 50,000 people or more, 1.7 percent of the units, carry 49.5 percent of the population, while the 697 settlements below 2,000 carry 12.3 percent. Any population-weighted statistic here is close to a statement about sixteen large towns and any unweighted statistic is close to a statement about several hundred villages. The two point in different directions, and Section 5 keeps them separate.

Table 2. Settlements analyzed beyond the two state capitals, by state and by population band.

	Settlements		Population, 2020	
Stratum	*n*	%	Persons	%
Osun	534	55.7	3,739,409	72.2
Kwara	424	44.3	1,439,930	27.8
500–2,000	697	72.8	638,816	12.3
2,000–10,000	186	19.4	795,575	15.4
10,000–50,000	59	6.2	1,179,177	22.8
50,000+	16	1.7	2,565,771	49.5
All	958	100.0	5,179,339	100.0

Note. A settlement is a contiguous cluster of populated WorldPop cells carrying at least 500 persons. Clustering uses 8-connectivity after a one-cell binary closing, which bridges carriageway and compound-wall gaps without merging genuinely separate places; population is summed over populated cells only, so the closing adds extent and never adds people. The 500-person floor is an analyst-set choice with a large consequence. The two states hold 8,334,767 persons and 7,099,585 of them, 85.2 percent, fall inside a qualifying cluster; the missing 14.8 percent is dispersed rural population living in farmsteads and single compounds, and it is excluded because neither an operator nor a road authority treats such a place as a unit to be served. That population is very unlikely to be better served than the villages that do qualify, so every exclusion figure in this paper is computed on the more favorable 85.2 percent. Exclusion rule for the capitals: a cluster goes if its polygon contains the OpenStreetMap place=city node for Osogbo or for Ilorin, which removes each capital's whole continuously built-up conurbation and nothing else. Two clusters qualify, Ìlọrin (Kwara), 1,137,723 persons over 278 km^2, and Òṣogbo (Osun), 782,523 persons over 174 km^2. The table reports what remains.

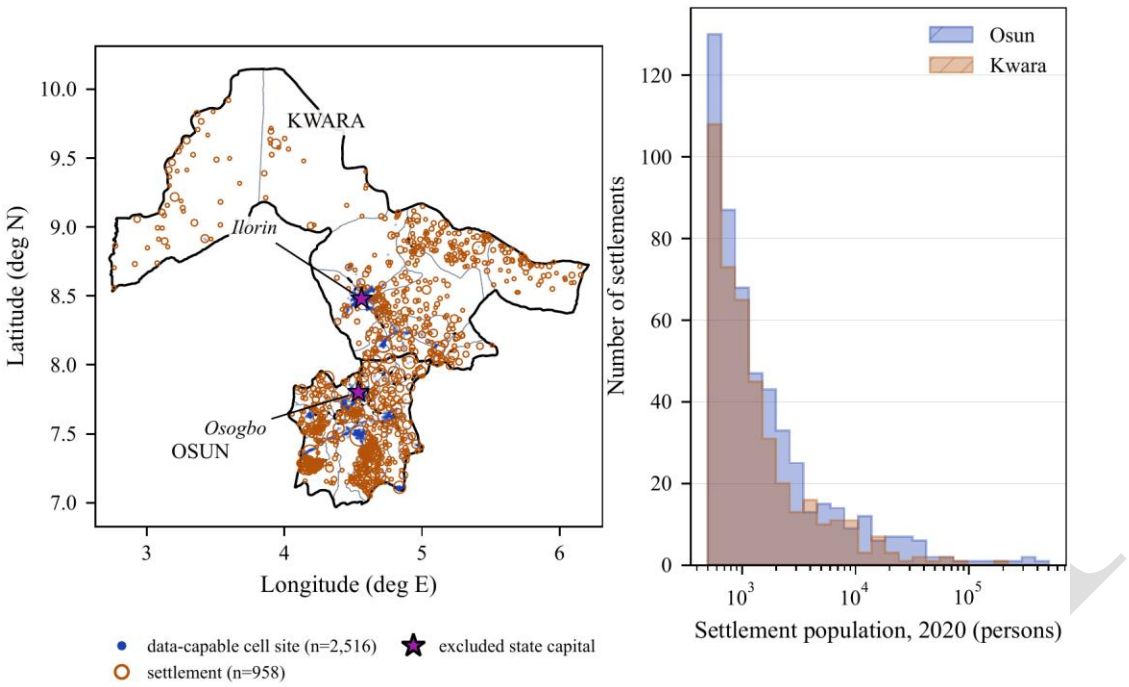

Figure 1. The study area, in two panels. Left, Osun and Kwara States over their local government boundaries, with the 958 analyzed settlements drawn as open circles scaled by population, the 2,516 data-capable cell sites lying inside the two state polygons as blue points, and the two excluded capitals starred. The panel is clipped to the two states, so masts in the wider retrieval box that shorten a border settlement's distance enter the computation without being drawn. Right, the settlement size distribution on a logarithmic population axis, split by state and separated by hatching as well as color, which shows the long right tail that puts half the study population into sixteen units. The exclusion rule is worth one sentence because it caused a real error. Osogbo and Ilorin are identified by the OpenStreetMap place=city node falling inside the cluster. Matching on the name string requires stripping Yoruba tone marks; a plain string match silently excluded neither capital on our first run, which would have put roughly a million well-served urban residents into a study explicitly about what lies beyond them.

4 Methods

Two properties of the condition layer belong at the head of this section. Condition is imputed for 53.3 percent of network length, by a class-conditional draw in the central case and by four fixed assumptions and one recalibration in the bounding scenarios. And the two tag families that can grade the same road independently agree at $\kappa = 0.041$, so the condition layer is unreliable at the level of an individual settlement and everything below is reported in aggregate for that reason. Section 6 gives the agreement measurement in full.

4.1 The five-point condition scale

Condition is scored on a five-point scale: grade 1 sealed, grade 2 engineered unsealed, grade 3 graded earth, grade 4 ungraded earth, grade 5 unmade. Grade 3 is the pivot of the scale and the source of most of its trouble, because surface=unpaved is a single undifferentiated OpenStreetMap value that lands there whatever the road is actually like.

The mapping from tag value to grade is fixed in advance and applied mechanically. From surface: asphalt, concrete and its variants, paving stones, sett, cobblestone, chipseal, bricks,

metal, wood and the generic paved give grade 1; compacted, fine gravel, gravel, pebblestone, grass paver, limestone, rock and stone give grade 2; unpaved, dirt, earth, ground, soil, clay and laterite give grade 3; sand, mud, grass and woodchips give grade 4. From tracktype, grade1 through grade5 map to grades 1 through 5 directly. From smoothness: excellent and good give grade 1, intermediate grade 2, bad grade 3, very bad grade 4, and horrible, very horrible and impassable grade 5. Where a way carries more than one tag, surface is read first, tracktype is used where surface is absent or gives only the generic grade 3, and smoothness is the last resort. The mapping is a judgment about what each tag value implies for a saloon car.

4.2 The two tests and the four classes

Each settlement is scored on two binary tests. It passes the coverage test if a data-capable operator site lies within the service radius of its built-up edge. It passes the condition test if at least 0.25 of the motorable centerline length lying inside its built-up extent plus a 1 km buffer, a rider's walk to a pickup point, sits at grade 1 or 2. The condition test is therefore a statement about the local network a vehicle would drive within the settlement and its approach, not about a single link to the classified network, and it is a share rather than a threshold on one road. Twenty-two of the 958 settlements have no mapped motorable road inside that buffer at all, which is itself an OpenStreetMap completeness finding; they are assigned their own state's network mean grade and flagged in the results file.

Crossing the two tests gives four classes: served on both, failing coverage only, failing condition only, and failing both. Population is attached to each class from the WorldPop cells composing the settlement. The quantity of interest is the union of the two failures, which is everything except the first class, because either failure alone disqualifies a settlement from a digitally booked ride. Separately from the binary tests, each layer also carries a continuous index used for the correlation and for the bivariate map. The coverage index is $0.7 \times proximity + 0.3 \times redundancy$, where proximity falls linearly from 1 at zero distance to 0 at twice the base radius and redundancy counts operators with a data-capable site inside the base radius, capped at two. The condition index is the rescaled mean grade. The weights 0.7 and 0.3, the cap of two operators and the linear rescaling are analyst-set. The correlation null below is computed on that composite; the constraint classes and the union rest on the two binary tests alone.

4.3 Testing whether the layers are correlated

If the two layers were positively correlated, the same settlements would tend to fail both tests and the union would be smaller than independence predicts. We measure the correlation four ways, since the answer could depend on how settlements are weighted: Pearson and Spearman on the continuous scores, Pearson on population-weighted scores, and the ϕ coefficient on the two binaries. We then compute the union directly and compare it against the union that independence would predict, which for marginal failure rates p_C and p_S is $p_C + p_S - p_C p_S$. The difference between observed and predicted is the quantity the hypothesis is about, and its confidence interval comes from a permutation test that shuffles one layer against the other across settlements.

4.4 The comparison that does not depend on the correlation

Separately we compare the union against each single margin. This asks a different and more operational question: how much larger is the excluded population than the worse of the two numbers a ministry would quote from its own sector. That comparison does not depend on the

correlation at all, which is why it survives when the correlation result does not.

4.5 The six surface-imputation scenarios

The condition layer is imputed for 53.3 percent of network length. We run six scenarios: a central class-conditional draw, four fixed assumptions treating every untagged link as sealed, made, ungraded earth or unmade, and a recalibration that re-reads every surface-derived grade through the empirical smoothness distribution for that grade. Section 6 carries all six. The headline figures quoted in the abstract and in the discussion are the central case; the bounds are reported alongside them wherever the central case appears in the text, and the reader should treat the band from 1.47 to 4.66 million as the honest width of the union estimate.

5 Results

5.1 The two layers separately

Table 3 and Figure 2 report the marginal distributions.

On coverage, 69.9 percent of the whole-state population falls within the service radius of a data-capable site, against the 74 percent third-generation population coverage the regulator reported for the fourth quarter of 2019 (Federal Ministry of Communications and Digital Economy 2020), a figure the international series reproduces to within a percentage point (International Telecommunication Union 2021). Second-generation coverage stood at 89 percent on the same basis, which our any-radio layer reaches at a 10 km radius and undershoots by 9.7 points at 5 km, and fourth-generation coverage stood at 37 percent, which this study cannot test at all because its extract holds twelve LTE records. Figure 2 draws the data-capable and any-radio cumulative curves together and marks the vertical gap between them at 5 km, which is the population that can place a call and cannot open a booking application.

On condition, the sealed share of network length is 17.6 percent in the central case, which includes imputed grades. The published national comparators do not agree with each other. The same World Bank assessment carries both a statement that 10 to 15 percent of the network is paved and a data box giving 20.3 percent for 2014 (World Bank 2019c), while a later appraisal puts it just above 30 percent, or 60,000 km, on a 2017 inventory (World Bank 2020). Our 17.6 percent sits above the first and just inside the second. Against the third, the most recent and therefore the closest to this paper’s period, it is low by a factor of 1.7. Counting only explicitly tagged sealed length, with the whole network still in the denominator so the comparison is like for like, the share falls to 9.3 percent, Osun 14.3 and Kwara 5.4. That is the floor from explicitly tagged length alone, 17.6 percent is the central estimate, and the published range is 10 to 30 percent. The coverage margin is therefore consistent with the published statistic for the layer this paper uses, and the condition margin is consistent with the older paved-share figures and not with the newest. That is one reason the condition layer is carried as a band from 1.47 to 4.66 million people rather than as a point.

Table 3. Marginal distribution of each constraint layer over the 958 settlements, unweighted and weighted by population.

	Passes the test		Distribution	
Layer and test	Settlements (%)	Population (%)	Median	Pop. w. median
Coverage: distance to the nearest data-capable cell site				
within 5 km	35.5	73.1	7.06 km	0.00 km
Condition: share of the local motorable network at grade 1 or 2				

at least 0.25	41.1	65.4	0.137	—
Association between the two layers				
Pearson r, indices	0.051 [-0.002, 0.110]			
Spearman ρ, indices	0.156 [0.100, 0.216]			
Pearson r, pop. weighted	-0.089 [-0.182, 0.028]			
Spearman ρ, pop. weighted	-0.050 [-0.182, 0.098]			
ϕ, the two binary tests	0.147 [0.083, 0.212]			

Note. $n = 958$ settlements carrying 5,179,339 persons. Distances are measured from the built-up edge, not the centroid. Intervals in square brackets are the 2.5th and 97.5th percentiles of 2,000 nonparametric bootstrap resamples of settlements. The coverage index combines proximity and the number of operators present on analyst-set weights of 0.7 and 0.3; the condition index is the rescaled mean grade. Both indices run from 0, worst, to 1, best, so a positive correlation means the two layers agree. By state, the unweighted Pearson r is 0.133 in Kwara and -0.022 in Osun. Every association in the lower block is computed on a condition layer whose two independent readings agree at $\kappa = 0.041$, so each estimate is attenuated toward zero by an unknown amount. Every point estimate lies between -0.09 and 0.16 and the sign is not stable across weighting or across the two states.

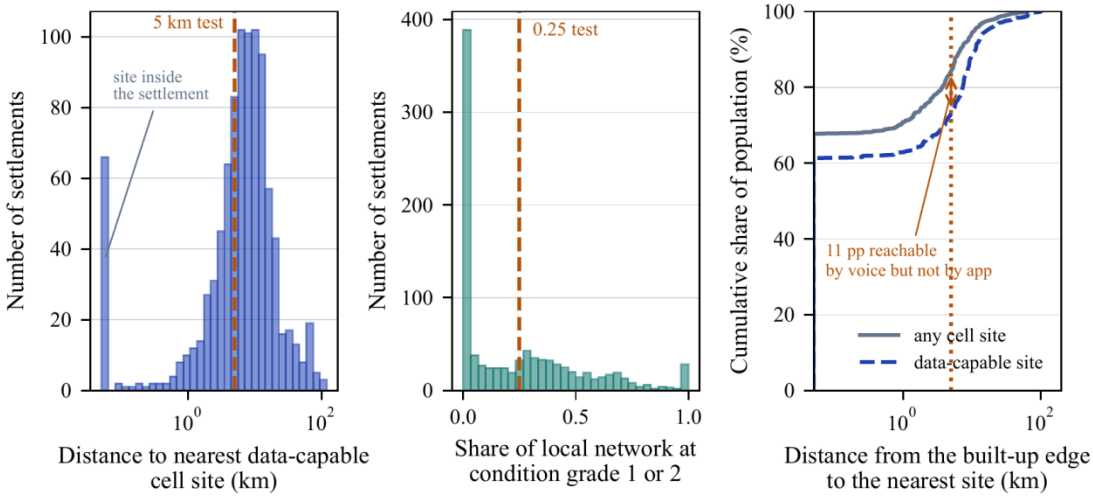

Figure 2. Each constraint layer on its own, in three panels. Left, the distribution of settlement distance to the nearest data-capable cell site on a logarithmic axis, with the 5 km test drawn and the settlements holding a site inside their own built-up extent marked at the left edge. Center, the distribution of the local made-surface share with the 0.25 test drawn, which shows how much of the frame sits just below the threshold. Right, cumulative population curves against radius for a site of any radio and for a data-capable site, with the vertical gap between them at 5 km marked. Across the 958 study settlements that gap is 11.3 percentage points, or 586,496 people, who are reachable by telephone and not by application; on the whole-state population including the capitals it is 9.4 points.

5.2 The correlation hypothesis fails

Every measure of association is small. Pearson correlation across settlements is 0.051 with a confidence interval of -0.002 to 0.110, Spearman is 0.156 with an interval of 0.100 to 0.216, the population-weighted Pearson is -0.089 with an interval spanning zero, and the ϕ coefficient on the binaries is 0.147 with an interval of 0.083 to 0.212 that excludes zero. The sign flips between weightings and between the two states, at 0.133 in Kwara and -0.022 in Osun.

These numbers cannot be read as evidence that the two layers are unrelated. The condition score entering every one of them is measured with error large enough that the layer disagrees with itself at $\kappa = 0.041$, and classical measurement error attenuates a correlation toward zero, so a small correlation is what a layer of this reliability produces almost regardless of the truth. The observed 0.051 is a lower bound on the magnitude of whatever association exists rather than an estimate of it, and we cannot bound the attenuation from above, because that needs a reliability coefficient these data do not supply: the 155 doubly tagged ways are a small and selected sample, the two tag families do not rate the same construct, and $\kappa = 0.041$ is itself a lower bound. What the null does establish is narrower and still useful, that nothing here licenses treating one layer as a proxy for the other.

Figure 3 maps the two indices against each other in terciles, both axes running in the direction of constraint severity so the worst-on-both class is the darkest corner of the key. If the layers tracked one another the map would show a diagonal, with the severe-on-both and easy-on-both classes each holding their own region and the off-diagonal classes rare. It does not: cells of every combination appear in both states, interleaved rather than nested. Figure 4 makes the same point on the scatter, with the two thresholds drawn as a frontier dividing the plane into the four constraint classes. Within each vertical strip of coverage distance, settlements occupy the full range of made share, and all four quadrants are populated. The one structure visible in either exhibit is size, not covariance.

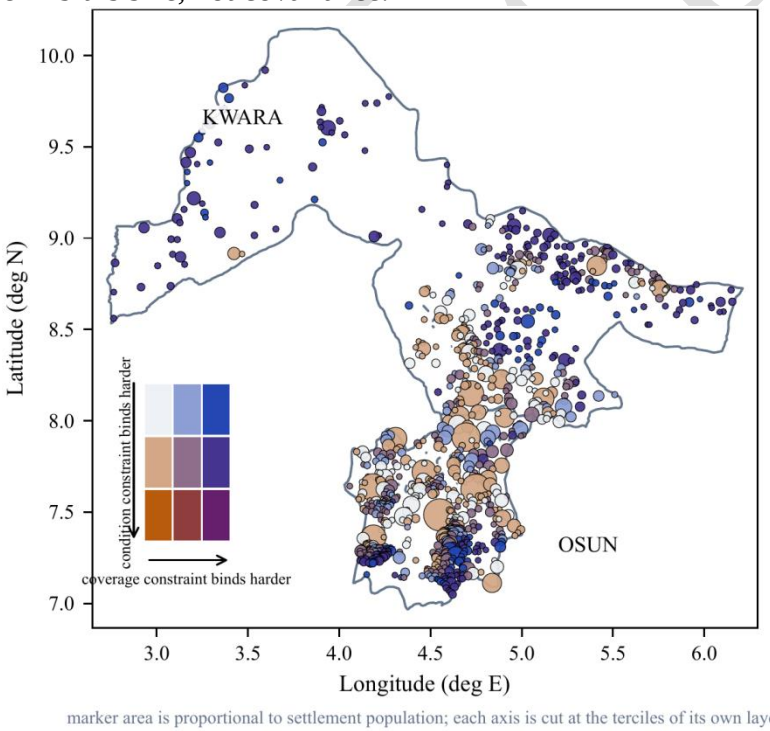

Figure 3. Bivariate choropleth of the two constraints across the two states, terciles on each axis, with the three-by-three key inset. Both axes run in the direction of constraint severity, so a

settlement in the darkest class is in the worst third on coverage and the worst third on condition simultaneously. The absence of a diagonal concentration is the correlation null in map form. The figure is a near-square map and is reproduced at its natural width rather than stretched to the text block.

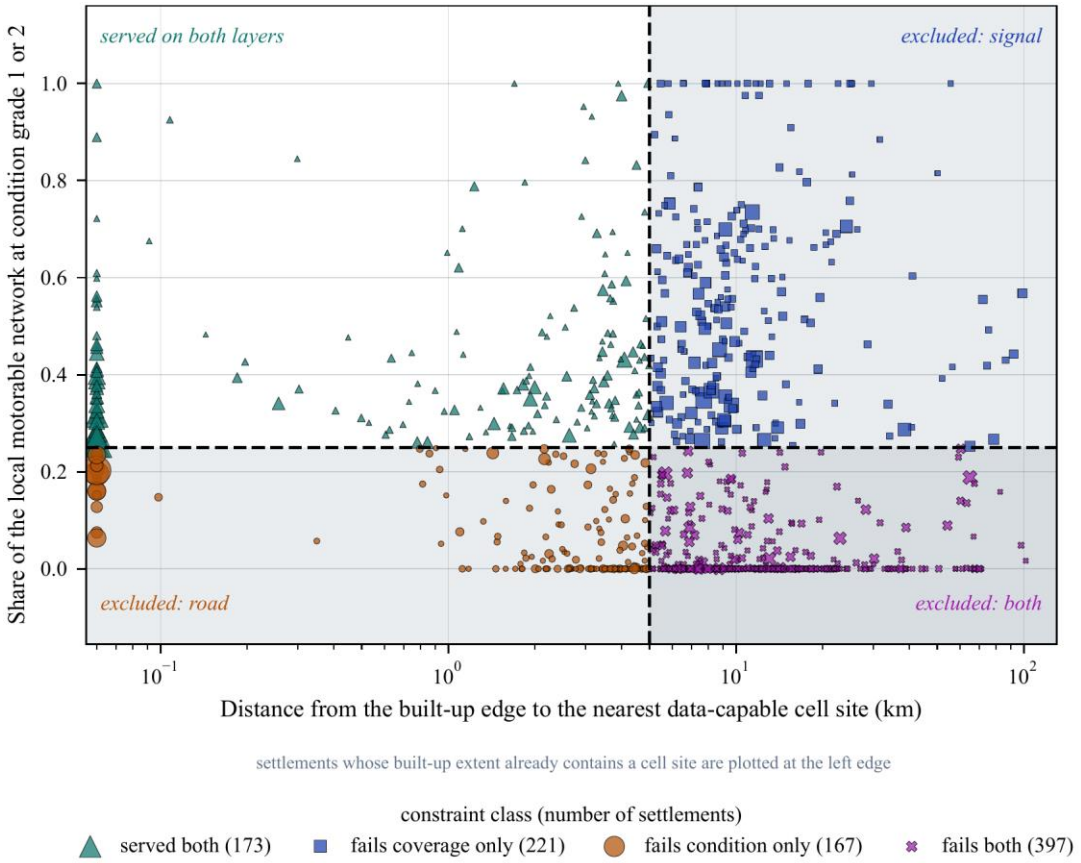

Figure 4. Coverage against condition for every settlement, with the 5 km and 0.25 thresholds drawn as a frontier and the four resulting quadrants named. Points carry the class color, marker and hatch used in every other exhibit. Settlements failing one test are spread across the full range of the other, which is the correlation null seen point by point rather than in aggregate.

The consequence for the union is exactly what independence predicts. Table 4 reports it. The observed excluded population is 2,538,738, or 49.02 percent of the study population, against an independence prediction of 2,700,527, or 52.14 percent. The difference is $-161{,}790$ people, or -3.12 percentage points, with a confidence interval running from -6.97 to $+0.98$ and a permutation test returning $p = 0.487$.

We set out to show that the correlation between these layers changes the size of the excluded population. It does not, in this study area, at this scale. Section 6 shows that even the sign of the difference is not stable.

Two components of uncertainty sit behind that permutation test and they are reported separately, because pooling them would produce a single interval that describes neither. **The resampling component** is the replicate standard deviation of the permutation distribution, which is 0.0737 over both states pooled and, computed within each state so the two are never mixed, 0.0889 in Osun and 0.0808 in Kwara. **The area component** is the difference between the two states' observed union shares, 39.5 percent in Osun against 73.6 percent in Kwara, a range of 34.1 percentage points. That is four times the resampling component and it is the larger source of

uncertainty about what this method would return in a third state. Two states is not enough for a standard deviation and none is quoted.

Table 4. The excluded population under the observed joint distribution of the two layers, against what statistical independence between them would predict.

Quantity	Both states	Osun	Kwara
Fails the coverage test	1,391,581	822,502	569,079
Fails the condition test	1,789,839	971,310	818,529
Union, observed	2,538,738	1,478,550	1,060,188
Union, under independence	2,700,527	1,580,168	1,064,114
Observed minus independence	-161,790	-101,618	-3,926
Union as a share (%)	49.02	39.54	73.63
Independence share (%)	52.14	42.26	73.90
Difference (pp)	-3.12	-2.72	-0.27
Union above the worse margin (%)	14.46	—	—
Union above the worse margin (persons)	+748,898	—	—

Note. pp, percentage points. Under independence the union share is $p_A + p_B - p_A p_B$, where p_A and p_B are the population shares failing each test. The observed union is 49.02 percent [34.27, 65.06] against an independence prediction of 52.14 percent, a difference of -3.12 pp $[-6.97, +0.98]$, which straddles zero. The permutation test shuffles the condition label among settlements within population quintiles, preserving the condition-to-size relationship and breaking only the coverage-to-condition association; over 2,000 draws it returns a null union of 54.58 percent [41.14, 68.86] and $p = 0.487$ two-sided. **Resampling and area variance are reported separately and never pooled.** The resampling component is the permutation replicate standard deviation, 0.0737 pooled, 0.0889 in Osun and 0.0808 in Kwara. The area component is the difference between the two states' observed union shares, 39.5 percent against 73.6 percent, a range of 34.1 pp; two states is not enough for a standard deviation and none is quoted. The last two rows give the quantity that survives the sweeps: the union exceeds the worse of the two single-sector statistics by 748,898 persons [454,655, 1,029,686], which is 14.46 pp [8.78, 19.88] at these thresholds and not at all thresholds, as Section 6 shows. Persons are rounded from full-precision population sums, so a difference taken between two rounded entries can sit one person away from the rounded difference. Every figure in this table rests on a condition layer measured at $\kappa = 0.041$.

5.3 Who is excluded, and by which failure

Table 5, Figure 5 and Figure 6 report the four classes.

Just over half the population beyond the capitals, 2,640,602 people or 51.0 percent, passes both tests. Of the remainder, 748,898 people or 14.5 percent fail on coverage alone, 1,147,157 or 22.1 percent fail on road condition alone, and 642,682 or 12.4 percent fail on both.

The shape of that breakdown is the finding. The single largest excluded group fails only on road condition, and would be invisible to any connectivity statistic. The second largest fails only on coverage, and would be invisible to any road statistic. Only 12.4 percent of the population, a quarter of the excluded, fails both and would appear in both sectors' numbers.

That separation is a population-weighted result and it does not hold when settlements are counted. Of the 785 excluded settlements, 397, or 50.6 percent, fail both tests, and the ϕ coefficient on the two binaries is 0.147 with a bootstrap interval of 0.083 to 0.212 that excludes zero. Counted as places, the two failures overlap heavily. Only when population is attached does the overlap shrink to a quarter of the excluded, and the reason is size structure rather than geography. Exclusion falls monotonically with settlement size, from 86.5 percent of the population below 2,000 people, through 70.7 and 50.9 percent, to 32.1 percent at 50,000 and above. The sixteen settlements of 50,000 or more carry 49.5 percent of the study population and all sixteen pass the coverage test; not one appears in the fails-coverage-only or fails-both classes. The population-weighted picture is therefore built around large towns that have signal and, in six cases, poor local roads, while the settlements failing both tests are overwhelmingly small. Both readings are true and they answer different questions. A ministry allocating capital across places should read the settlement counts; one counting people should read the population columns and should know that the separation it sees there is a property of the size distribution.

Table 5. Settlements and population in each of the four constraint classes, for the two states together and separately.

	Both states		Osun		Kwara	
Constraint class	n	Persons	n	Persons	n	Persons
Served on both layers	173	2,640,602	112	2,260,860	61	379,742
Fails coverage only	221	748,898	131	507,239	90	241,659
Fails condition only	167	1,147,157	100	656,047	67	491,110
Fails both	397	642,682	191	315,263	206	327,419
Excluded, the union	785	2,538,738	422	1,478,550	363	1,060,188

Note. Base-case tests: a data-capable cell site within 5 km of the built-up edge, and at least 0.25 of the local motorable network at condition grade 1 or 2, the latter measured on a layer whose two readings agree at $\kappa = 0.041$. A settlement failing either test is excluded from digitally booked service, so the excluded population is the union of the last three rows, 2,538,738 persons or 49.0 percent of the 5,179,339 living beyond the two capitals. **The n columns and the person columns tell different stories and both are needed.** By persons, the two constraints bite in different places: 748,898 have a usable road and no data signal while 1,147,157 have a data signal and no usable road. By settlements, they bite in the same places: 397 of the 785 excluded units, 50.6 percent, fail both. The gap between the two readings is the settlement size distribution, since exclusion runs from 86.5 percent of population below 2,000 persons to 32.1 percent at 50,000 and above. Class populations are rounded from full-precision sums and need not add to the rounded union.

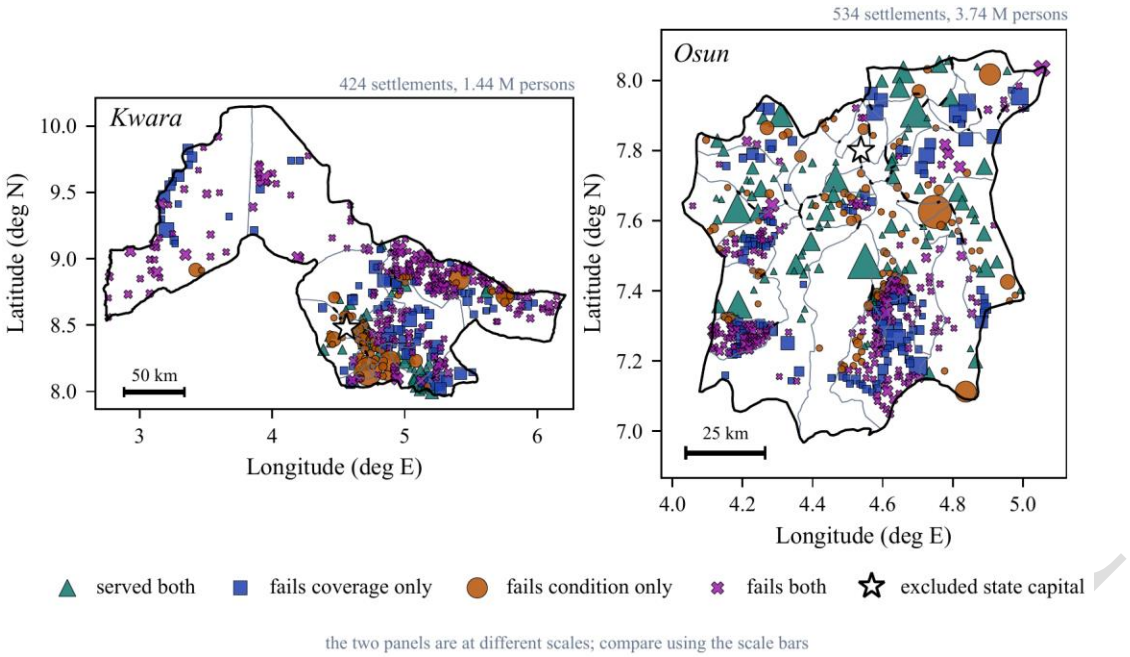

Figure 5. Constraint typology mapped, one panel per state at its own scale, with a scale bar on each and a note that the two scales differ. Each settlement takes the color, marker and hatch of its class: served on both, failing coverage only, failing condition only, failing both. The hatching carries the classes independently of color so the map survives greyscale printing. Read against Figure 3, the panel shows that the fails-both class is numerous and small while the served-both class is fewer and larger, which is the size structure discussed in the text.

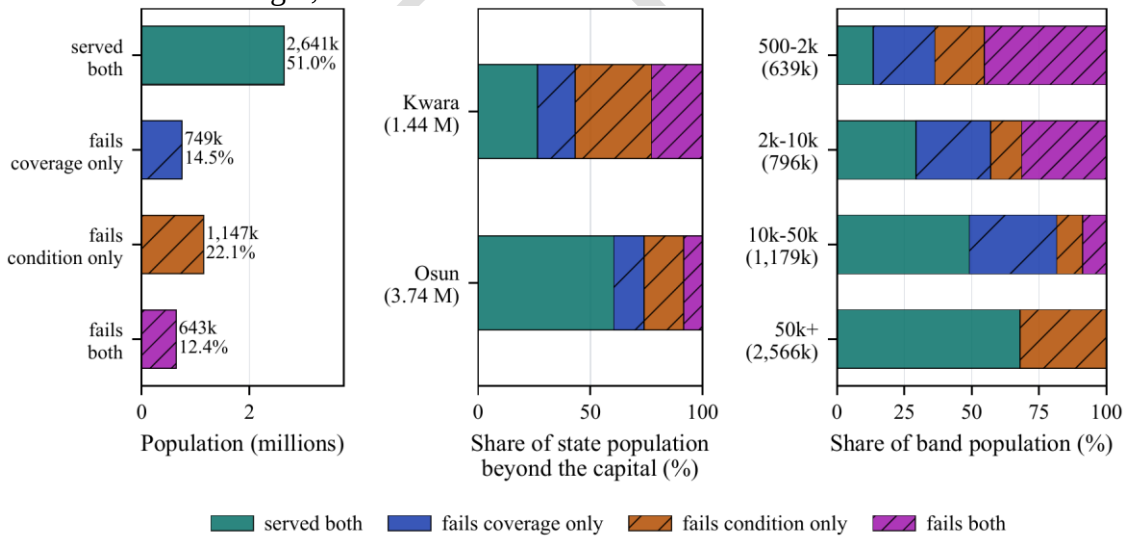

Figure 6. Population at risk by constraint class, in three panels: overall, by state, and by settlement size band. The overall panel shows that the largest excluded group fails only on road condition and is invisible to any connectivity statistic, while the second largest fails only on coverage and is invisible to any road statistic. The by-state panel shows Kwara at 73.6 percent excluded against Osun at 39.5, a 34.1-point spread that is the paper's area variance. The by-size panel shows the monotone gradient, from 86.5 percent of population excluded below 2,000 persons to 32.1 percent at 50,000 and above, with the fails-coverage-only and fails-both classes empty in the largest band.

5.4 The union exceeds either margin, robustly

Because the two failing populations overlap only partly by population, the union is larger than either margin alone. At the base thresholds it exceeds the worse of the two single-sector statistics by 748,898 people, which is 14.46 percentage points, with a confidence interval from 8.78 to 19.88. Unlike the correlation result, the sign of this comparison is close to universal across the threshold space, though its size is not. Figure 7 draws it across the 56 cells of the radius-by-threshold sweep. The union exceeds the worse margin in 55 of them and equals it exactly in the 56th, at a 1 km radius and a 0.05 made-surface threshold, where the coverage test is so strict that every settlement failing the condition test is already failing the coverage test. The margin ranges from 0.00 to 18.85 percentage points across the sweep, is 14.46 points at the base thresholds, and falls below 1.3 points in five further cells, all of them at the extremes of the grid. So the direction is stable and the magnitude is not: a reader is entitled to conclude that a single-sector statistic understates the excluded population, and is not entitled to conclude that the understatement is fourteen points anywhere but at these thresholds.

That is the operational claim of this paper, stated at its correct strength. A ministry quoting its own sector's worst number understates the population that cannot use a digitally booked service, by fourteen percentage points at the thresholds we adopt and by an amount that has to be recomputed at any other, and it does so not because its number is wrong but because the people its number misses are failing a test it does not administer.

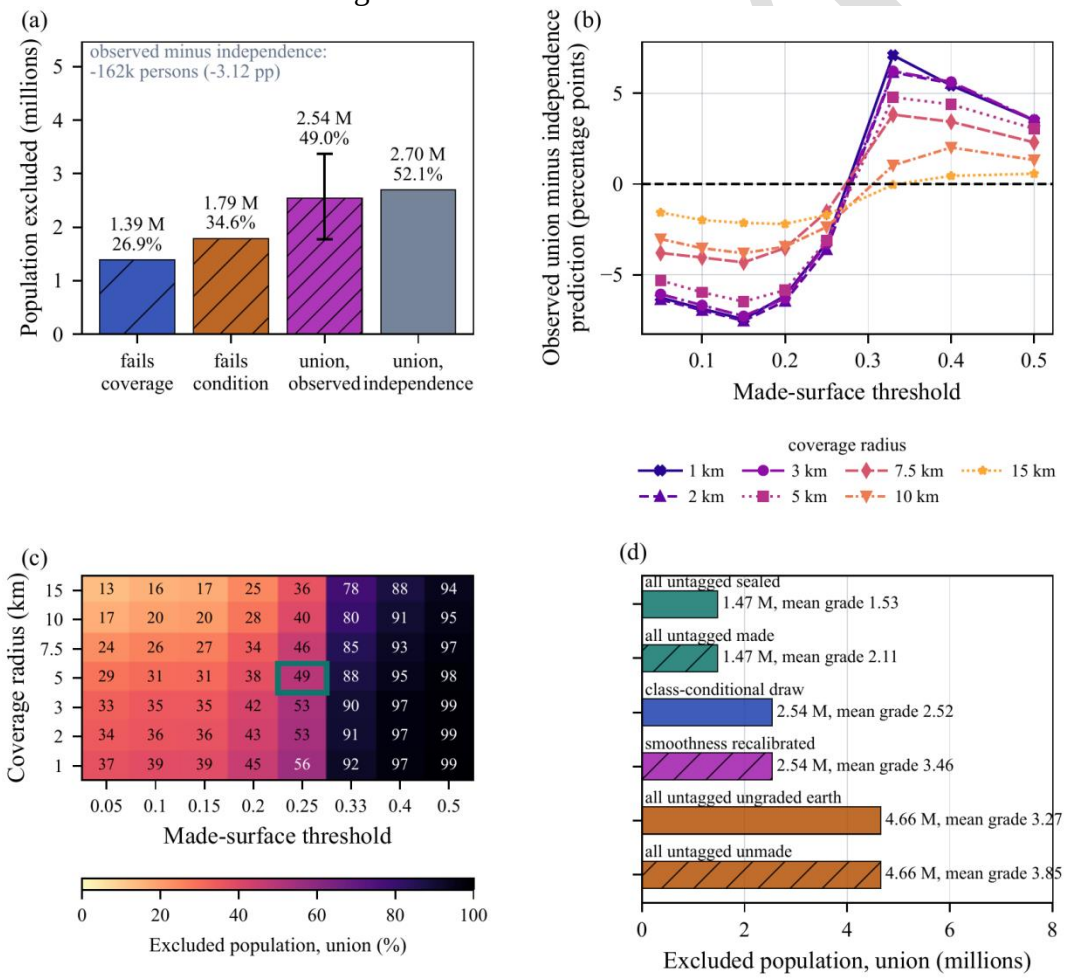

Figure 7. The union and what it rests on, in four panels. (a) The two single-sector margins, the

observed union and the independence prediction as populations, with the bootstrap interval on the observed union and the observed-minus-independence difference annotated. (b) That difference traced across the eight made-surface thresholds, one line per service radius, with the zero line drawn: the curves cross it near 0.30, which is the sign reversal the paper reports. (c) The full seven-by-eight grid of radii against thresholds, shaded and labeled by the union as a share of population, running from 13 to 99 percent, with the base cell at 5 km and 0.25 outlined; the union-over-worse-margin difference computed on the same 56 cells runs from 0.00 to 18.85 percentage points, positive in 55 of them and exactly zero at 1 km and 0.05. (d) The surface-imputation envelope, six scenarios from every untagged link sealed to every untagged link unmade, each labeled with its union and its population-weighted network mean grade, which is the 1.47 to 4.66 million band rather than the central case alone.

6 Sensitivity and Robustness

6.1 The threshold sweep

Table 6 and Figure 7 report seven service radii crossed with eight condition thresholds, 56 cells in all. The union itself ranges from 13.4 to 99.3 percent across that space, which is the expected consequence of moving two thresholds that jointly define exclusion, and is why we report the comparison against the margins rather than the level.

One region of that space deserves naming rather than smoothing. Between made-surface thresholds of 0.25 and 0.33 the population-weighted result moves sharply: the share failing the condition test rises from 34.6 percent to 76.5 percent and the union rises from 2,538,738 to 4,538,392 people. That happens because a small number of large towns sit with local made shares between the two values and cross the line together, which is the size structure of the frame reappearing as a sensitivity. Our base threshold of 0.25 sits at the foot of that cliff, and it was chosen before the cliff was found rather than because of where it sits. Panel B of Table 6 reports the 0.33 cell alongside its neighbors so a reader can see the jump rather than infer it.

Table 6. Excluded population across the threshold sweep and across every surface-imputation scenario.

Scenario	Network mean grade	Fails condition (%)	Union (persons)	Union above the worse margin (pp)
Panel A. Surface imputation, base thresholds held				
All untagged sealed (grade 1)	1.53	6.5	1,473,142	+1.57
All untagged made (grade 2)	2.11	6.5	1,473,142	+1.57
Class-conditional draw	2.52	34.6	2,538,738	+14.46
Smoothness recalibrated	3.46	34.6	2,538,738	+14.46
All untagged ungraded earth (grade 4)	3.27	80.1	4,656,153	+9.81
All untagged unmade (grade 5)	3.85	80.1	4,656,153	+9.81
Panel B. Threshold sweep, central imputation				
radius 1 km, made share 0.05	—	10.0	1,921,025	0.00
radius 1 km, made share 0.10	—	13.4	2,000,612	+1.54
radius 5 km, made share 0.10	—	13.4	1,588,869	+3.81
radius 5 km, made share 0.25	—	34.6	2,538,738	+14.46
radius 5 km, made share 0.33	—	76.5	4,538,392	+11.09

radius 5 km, made share 0.40	—	87.0	4,915,642	+7.87
radius 10 km, made share 0.25	—	34.6	2,047,907	+4.98
radius 15 km, made share 0.25	—	34.6	1,888,826	+1.91
radius 15 km, made share 0.50	—	93.4	4,883,949	+0.93
Panel C. Alternative condition test: distance to a sealed classed road				
radius 5 km, within 1 km	—	13.3	1,558,195	+3.22
radius 5 km, within 2 km	—	10.2	1,501,324	+2.12
radius 5 km, within 5 km	—	4.3	1,408,384	+0.32
radius 5 km, within 10 km	—	2.0	1,391,581	0.00

Note. pp, percentage points. The condition layer underlying every row agrees with itself at $\kappa = 0.041$, and Panel A reports the full width of the resulting band rather than a robustness check that passes. Panel A holds both base thresholds and varies only what is assumed about the 76,889 ways carrying no surface, tracktype or smoothness tag, which is 53.3 percent of routed length; there are six scenarios, not three. The made-surface test cannot distinguish grade 1 from grade 2, nor grade 4 from grade 5, so the four fixed scenarios collapse to two union values, 1,473,142 and 4,656,153 persons; the network mean grade column shows that the underlying condition surfaces do differ. **The smoothness-recalibrated row repeats the central row digit for digit, and that is arithmetic rather than a copying error.** The recalibration re-reads only grades that came from a surface value, and among the 155 doubly tagged ways only the surface grade 3 reading carries enough observations to supply a conditional distribution; that distribution puts all of its mass on grades 3, 4 and 5. No length crosses into or out of grades 1 and 2, so the made-surface share the condition test is defined on is identical settlement by settlement and every population figure that follows from it is identical too. The mean grade moves from 2.52 to 3.46 because it is computed over all five grades. Panel B reports nine of the 56 cells of the full seven-by-eight sweep, which is archived in full with the manuscript and drawn in Figure 7; across the whole sweep the union runs from 13.4 to 99.3 percent of the population, the observed-minus-independence gap runs from -7.55 to $+7.10$ pp and changes sign near a made-surface threshold of 0.30, and the union exceeds the worse single-sector margin in 55 of the 56 cells and equals it in the 56th, by between 0.00 and 18.85 pp. Panel C reports four of the 42 cells of a second sweep that replaces the local made-share test with distance from the built-up edge to the nearest sealed classed road; there the margin runs from 0.00 to 6.55 pp and is exactly zero in nine cells. The direction of the correlation effect is not robust; the sign of the margin comparison is nearly so; its magnitude is not. The observed-minus-independence gap runs from -7.55 to $+7.10$ percentage points and changes sign near a made-surface threshold of 0.30. Its direction is therefore not robust, which is a stronger statement of the null than the central estimate alone: not only is the difference small, but which way it points depends on a threshold we chose.

The union-exceeds-the-worse-margin comparison is positive in 55 of the 56 cells and zero in the 56th, which is why it rather than the correlation carries the paper. Its size ranges from 0.00 to 18.85 percentage points, so the direction survives the sweep and the magnitude does not.

The coverage layer is insensitive to processing choices. Changing the site deduplication grid by a factor of twenty moves the result by 0.11 percentage points.

6.2 A second sweep, on a different definition of adequate condition

The condition test above measures the share of a settlement's own local network that is made. That is one reasonable operationalization and not the only one, so we computed a second sweep in which the condition test is instead the distance from the built-up edge to the nearest sealed

classed road, crossing seven service radii with six distance thresholds for 42 cells. Panel C of Table 6 reports four of them and all 42 are archived with the manuscript. It weakens the paper's headline. The margin collapses: across the 42 cells the union exceeds the worse single-sector margin by between 0.00 and 6.55 percentage points, against 0.00 to 18.85 on the paper's own test, so the largest value anywhere in the second sweep is barely a third of the largest in the first. The margin is also exactly zero in nine of the 42 cells, 21 percent of them, all at the loosest distance thresholds, where every settlement failing the condition test already fails the coverage test and the union is just the coverage margin. And the association between the layers strengthens: ϕ runs from 0.089 to 0.413 across the 42 cells, against the base-case 0.147 the paper elsewhere calls weak, while the observed union sits below the independence prediction in every cell, by between 1.23 and 8.10 points. The reason is mechanical and it constrains how the headline should be read. A settlement far from any sealed classed road is also, in this region, far from the towns where masts are built, so the two failures line up. A settlement with a poor local street network may sit beside a sealed highway, so they do not. The margin the paper reports is therefore a property of a condition test defined on the local network, which we take to be the right test for a ride that has to reach a doorstep, and not a property of road access in general. The correction is real in direction under both definitions and its size is not portable between them.

6.3 The imputation bounds

The excluded population runs from 1.47 million if every untagged link is treated as made to 4.66 million if every untagged link is unmade, against a central case of 2.54 million. That is a wide band. What is less sensitive to it is the comparison against the margins, because the imputation moves both the condition margin and the union together: the margin is $+1.57$ points at the optimistic bound, $+14.46$ at the central case and $+9.81$ at the pessimistic one, all positive and none close to each other. A recalibration scenario using the smoothness tag to reset grades worsens the mean condition grade from 2.52 to 3.46 and leaves every population figure unchanged. It moves length only among grades 3, 4 and 5 and never into or out of grades 1 and 2, so the made-surface share the condition test reads is the same length in both scenarios and the two rows of Panel A agree exactly.

The central case is not the middle of the evidence. Three independent facts push the same way. The condition layer's two readings differ by a mean of -1.15 grades, with the surface reading, the one the layer mostly uses, systematically the more optimistic. Our rural access comparator overshoots the published index by 7.9 points in Osun and 13.4 in Kwara, in both cases because we can apply no condition filter and the World Bank's can. And the World Bank's own condition evidence for Nigeria is severe: 40 percent of federal, 78 percent of state and 87 percent of local government roads in poor condition (World Bank 2019b), and under 2 percent of the unpaved network in good condition (World Bank 2019c). **If that evidence is taken at face value, the true union for these two states sits nearer the pessimistic bound of 4.66 million people, 89.9 percent of the population beyond the capitals, then the central case of 2.54 million.** We continue to headline the central case, because it is the one the observed tags support without importing a national condition statistic into a settlement-level layer, but a reader deciding how much to trust 2.54 million should treat it as a floor rather than a midpoint.

6.4 What the published benchmarks return

Coverage at 69.9 percent sits close to the 74 percent third-generation figure for the period, the sealed share is consistent with the two older published paved-share figures though not with the

newest, and the ordering of the two states matches their published rural access indices, which are 57.3 percent in Osun and 21.7 in Kwara against a national 25.5 (World Bank 2019c). Those checks pass. The contrast between the two states is unusually wide: Osun records the highest state rural access index in Nigeria and Kwara sits below the national figure, so the two states in this study bracket most of the national range.

Two fail. Our rural access comparator overshoots the published index by 7.9 percentage points in Osun and 13.4 in Kwara, which is why the pessimistic imputation bound is promoted to a reported scenario rather than left in an appendix. And the any-radio layer understates second-generation reach by 9.7 percentage points at the 5 km radius, which is why both layers, data-capable and any-radio, are reported side by side above rather than collapsed into a single coverage number. A further check, on state-level fund allocations, was demoted from a benchmark to a limitation because the published clusters are not comparable to our frame.

6.5 What "rural" means in the access figure

We report 65.2 percent for the rural population of Osun within 2 km of a classed road carrying an explicitly tagged sealed surface, against the published 57.3. A figure of this kind is far more sensitive to how "rural" is defined than to anything in the road computation, so the definition is worth stating on its own.

The road side is straightforward and does little of the work. We restrict to trunk, primary, secondary and tertiary, keep only ways whose sealed surface is explicitly tagged rather than imputed, densify those ways to points and count population within 2 km. Computed on the whole state population rather than on the rural part alone, the same construction returns 77.2 percent.

The denominator is where the sensitivity sits. This paper classifies cells, calling a cell rural when it lies outside any settlement cluster of 50,000 persons or more. The common alternative is to classify whole local government areas, calling an area rural when its density falls below 1,000 persons per square kilometer and counting everybody living in one. The cell rule is the stricter of the two, because it removes the built-up core of every large town wherever it sits, including towns inside low-density local government areas that the area rule keeps, and those cores are exactly the places near a tagged sealed classed road. Applied to the same road layer, the area rule would therefore return a materially higher figure than 65.2 percent, on the same distance computation and the same tags.

For comparison against a rural access index, which is defined on rural population rather than on rural administrative units, the cell-based rule is the closer match, so 65.2 percent is the figure we carry forward. It still overshoots the published 57.3 for the reasons already given, and the overshoot is smaller than the looser rule would produce.

6.6 Reliability of the condition layer

The condition layer is scored by the tag mapping set out above, and its reliability has to be established on the data the layer is actually built from. A Cohen’s κ of 0.561 between two survey teams grading a road inventory in the field would be a different measurement on a different object: it says nothing about how consistently two OpenStreetMap tag families grade the same road, and the conditions under which such a figure is obtained do not carry over.

We measured agreement directly on the tagged sample instead, using the 155 OpenStreetMap ways that carry both a surface and a smoothness tag and can therefore be graded twice independently. Exact agreement is 18.1 percent and within-one-grade agreement is 61.9 percent. Cohen’s κ (Cohen 1960) is 0.041, with a bootstrap interval from 0.0002 to 0.083. Because the

scale is ordered and a two-grade disagreement is worse than a one-grade disagreement, we also report the weighted forms: linearly weighted κ is 0.071 with an interval from -0.013 to 0.154, and quadratically weighted κ is 0.098 with an interval from -0.048 to 0.230. Both weighted intervals straddle zero, so weighting disagreements by their size does not rescue the layer.

The cross-tabulation shows the mechanism. Of the 155 ways, 137 read as grade 3 from surface, because surface=unpaved is a single undifferentiated value, while their smoothness values spread across grades 3, 4 and 5. The surface tag records what a road is made of and the smoothness tag records what condition it is in, and those are not the same question. That also means the two are not two observers rating the same thing, so the figure is a lower bound rather than a clean reliability estimate. It is close to no agreement at all, it is the most important limitation in this paper, and it is symptomatic rather than idiosyncratic. The published road condition record for these two states is close to empty: the national statistics office publishes no road condition data, the federal maintenance agency publishes none on its own site, and the best available national split is a single assessment reporting 40 percent of federal, 78 percent of state and 87 percent of local government roads in poor condition (World Bank 2019b). A layer built from volunteered tags disagrees with itself because there is nothing better to build it from.

The error is not only large, it has a direction, and a direction does not average out. The mean signed difference between the two readings is -1.15 grades, with the surface reading, which supplies most of the layer, systematically the more optimistic. Aggregation cancels the random part of a disagreement and leaves the systematic part untouched, so the honest statement is not that the grading error washes out at population level but that the aggregate condition figures here are biased optimistic by roughly one grade. That bias runs in the same direction as the rural access overshoot above and as the gap between our sealed share and the newest published paved share. We report at aggregate level only, because settlement-level noise makes anything finer meaningless, and those aggregates should be read as a lower bound on exclusion rather than as unbiased estimates.

7 Discussion

7.1 Interpreting the near-zero association

The correlation hypothesis was the reason we built this study, and it failed. It is reported in full, because the sweep that establishes the null also establishes what survives.

The null does not convert into a claim about investment history. It would be easy to write that connectivity investment and road investment in this region have not tracked one another, and that this is a finding about how the two sectors developed. That sentence does not follow from these data. The correlation is computed on a condition layer that disagrees with itself at $\kappa = 0.041$, and classical measurement error in one of two correlated variables attenuates the observed correlation toward zero in proportion to the layer's unreliability. A near zero correlation is what a layer this noisy produce whether the underlying association is zero or substantial. To bound the attenuation, we would need a reliability coefficient for the condition score, and the only agreement statistic available is itself a lower bound computed on 155 purposively double-tagged ways whose two readings measure different constructs. We cannot bound it. The null is a statement about what these data can support, not about how Nigerian ministries have spent money. Two weaker statements do survive, and they are the ones the rest of the paper uses. Nothing in these data licenses treating either layer as a proxy for the other, which is enough to justify building the joint distribution rather than inferring it. And the alternative condition test in

Section 6, which is not obviously worse than the one we adopt, returns ϕ values as high as 0.413 and an observed union below the independence prediction in every one of its 42 cells, which is a further reason to treat the near-zero association as a property of a particular operationalization rather than a property of the region.

There is a practical consequence and it cuts against the paper's own convenience. An analyst holding the coverage statistic and the road statistic and no joint data cannot compute the excluded population. The independence assumption they would have to make happens to be approximately right here, but that is not knowable in advance, and its error runs in either direction depending on where the thresholds sit and which condition test is used. Everything that needs a cell of the joint table needs the joint table.

7.2 What follows for the two ministries

Three organizations hold the pieces of this problem and none holds all of them. The Federal Ministry of Communications and Digital Economy owns the coverage layer through the Universal Service Provision Fund, the Nigerian Communications Commission publishes the coverage statistics that the fund's targets are written against, and the Osun and Kwara state ministries of works own the roads. The recommendations below are addressed to them in that order. First, publish the joint distribution rather than the two margins. Both layers already exist in open data, the crossing is arithmetic, and the resulting number is the one that describes who can actually use a service. Neither ministry can produce it alone, which is precisely why it does not exist. Second, target the intersection deliberately if the objective is service access. Only 12.4 percent of the population fails both tests, so an intervention aimed at that group reaches a quarter of the excluded population and no more. Conversely, an intervention on either single layer leaves at least as many people excluded as it helps, because the other failure is still there. Note that the same group is 397 of the 785 excluded settlements, so an agency counting places rather than people faces a very different arithmetic and should say which it is counting.

Third, do not treat 14.46 percentage points as a portable correction factor. **The margin cannot be computed from the margins.** It is identically the population failing coverage and passing condition divided by the total, which is a cell of the joint table, so anyone holding only the two marginal statistics cannot recover it. That is the impossibility named above. Nor is it one number: it runs from 0.00 to 18.85 points across the threshold sweep and from 0.00 to 6.55 across the alternative condition sweep. What a ministry can safely take from this paper is the direction and the method. Its own sector's worst statistic understates the excluded population, the understatement is large enough to matter, and the only way to size it is to cross the two layers, which costs a day of work on data both bodies already hold.

7.3 Transferability

The method transfers wherever both layers exist in open data, which is now most places, and the frame requires no proprietary input. Neither the correlation null nor the size of the margin transfers, and neither should be assumed. In a region where both investments followed the same development gradient the layers would correlate, the union would be smaller than independence predicts, and the margin would be smaller. In a region whose condition data is better than ours the correlation would also be less attenuated, which could move the result in either direction. Both are empirical questions in each setting, and the point of this paper is that both can be answered cheaply once the joint table exists.

7.4 Limitations

The condition grading does not agree with itself, at $\kappa = 0.041$ on the only double-graded sample available, and the weighted forms at 0.071 and 0.098 carry intervals straddling zero. Aggregate results are defensible and settlement-level ones are not. The disagreement is also directional, at a mean signed difference of -1.15 grades with the reading the layer mostly uses the optimistic one, so aggregate condition figures here are biased toward understating exclusion by roughly a grade, and every correlation in this paper is attenuated by an amount we cannot bound.

The cell layer is old. Its median record was last updated in October 2016 and none was touched in 2021, so five years of mast building are missing from a layer used to describe 2021 service. That understates coverage and therefore inflates both the union and the margin. The bias runs opposite to the condition layer's, but the two do not cancel in any calculable way and neither is corrected.

8 Conclusion

We built the joint distribution of mobile coverage and road surface condition for 958 settlements and 5.18 million residents beyond the capitals of Osun and Kwara, entirely from open data, to ask whether the correlation between the two layers determines how many people a digitally booked service cannot reach.

It does not. Every measure of association is small, at a Pearson correlation of 0.051 and a permutation test returning $p = 0.487$, and the observed excluded population differs from the independence prediction by 3.12 percentage points in a direction that reverses under a modest change of threshold. The hypothesis the study was built to test fails, and it fails on a condition layer that disagrees with itself at $\kappa = 0.041$. That is not evidence the two sectors developed independently; it is evidence these data cannot answer the question.

What holds instead is a comparison against the margins, which never asks how the layers covary. The union of the two constraints exceeds the worse single-sector statistic by 748,898 people, or 14.46 percentage points, at the base thresholds. It exceeds it in 55 of the 56 threshold combinations examined and equals it in the 56th, by between 0.00 and 18.85 points, and under an alternative definition of adequate road condition by between 0.00 and 6.55 points. The direction is close to universal and the magnitude is local to a choice of test. Half the population beyond the capitals fails at least one test, and by population only a quarter of the excluded fail both, though by settlement count half of them do.

The recommendation is arithmetic rather than infrastructural. Both layers are already open, the crossing costs nothing, and neither ministry can perform it alone. Publishing the joint distribution would replace two statistics that each describe part of the problem with one that describes it.

The layer that limits this study is condition, and no amount of further processing will fix it. Two volunteered tags asked different questions of the same road and returned answers that agree at $\kappa = 0.041$; that is a ceiling on what OpenStreetMap can tell anyone about Nigerian road surfaces, not a defect in how we read it. Driving a few hundred approach roads with a clipboard and the five-point scale used here, in the settlements the typology already identifies as borderline, would produce the first real condition inventory these two states have and would settle in a fortnight what the whole imputation apparatus in Section 6 is built to work around.

References

Adeleke, Richard. 2021. "Digital Divide in Nigeria: The Role of Regional Differentials." *African Journal of Science, Technology, Innovation and Development* 13 (3): 333–46. https://doi.org/10.1080/20421338.2020.1748335.

Alcorn, Louis G., and Alex Karner. 2021. "Integrating Formal and Informal Transit into One Hybrid Passenger Transport System in Lagos, Nigeria." *Transportation* 48 (3): 1361–77. https://doi.org/10.1007/s11116-020-10099-8.

Alkire, Sabina, and James Foster. 2011. "Counting and Multidimensional Poverty Measurement." *Journal of Public Economics* 95 (7–8): 476–87. https://doi.org/10.1016/j.jpubeco.2010.11.006.

Alkire, Sabina, James E. Foster, Suman Seth, Maria Emma Santos, Jose Manuel Roche, and Paola Ballon. 2015. *Multidimensional Poverty Measurement and Analysis: Chapter 5 – the Alkire-Foster Counting Methodology*. OPHI Working Paper No. 86. Oxford Poverty and Human Development Initiative (OPHI). https://ophi.org.uk/sites/default/files/OPHIWP086_Ch5.pdf.

Anwar, Mohammad Amir, and Mark Graham. 2020. "Hidden Transcripts of the Gig Economy: Labour Agency and the New Art of Resistance Among African Gig Workers." *Environment and Planning A: Economy and Space* 52 (7): 1269–91. https://doi.org/10.1177/0308518X19894584.

Atkinson, A. B. 2003. "Multidimensional Deprivation: Contrasting Social Welfare and Counting Approaches." *The Journal of Economic Inequality* 1 (1): 51–65. https://doi.org/10.1023/A:1023903525276.

Bahia, Kalvin, Pau Castells, Genaro Cruz, et al. 2020. *The Welfare Effects of Mobile Broadband Internet: Evidence from Nigeria*. IZA Discussion Paper No. 13219. Institute of Labor Economics (IZA). https://docs.iza.org/dp13219.pdf.

Barrington-Leigh, Christopher, and Adam Millard-Ball. 2017. "The World's User-Generated Road Map Is More Than 80% Complete." *PLOS ONE* 12 (8): e0180698. https://doi.org/10.1371/journal.pone.0180698.

Barron, Christopher, Pascal Neis, and Alexander Zipf. 2014. "A Comprehensive Framework for Intrinsic OpenStreetMap Quality Analysis." *Transactions in GIS* 18 (6): 877–95. https://doi.org/10.1111/tgis.12073.

Bello, Olayiwola Wasiu, Jayeola Femi Opadiji, Nasir Faruk, and Yinusa Ademola Adediran. 2016. "Opportunities for Universal Telecommunication Access in Rural Communities: A Case Study of 15 Rural Villages in Nigeria's Kwara State." *The African Journal of Information and Communication*, no. 17: 139–63. https://ajic.wits.ac.za/article/view/13622.

Bianchini, Alessandra, Paola Bandini, and David W. Smith. 2010. "Interrater Reliability of Manual Pavement Distress Evaluations." *Journal of Transportation Engineering* 136 (2): 165–72. https://doi.org/10.1061/(ASCE)0733-947X(2010)136:2(165).

Bolt Technology OÜ. 2020. *Bolt in Ilorin*. Operator city page, Internet Archive capture of 30 December 2020, https://web.archive.org/web/20201230051507/https://bolt.eu/en-ng/cities/ilorin/.

Bolt Technology OÜ. 2021a. *Bolt in Osogbo*. Operator city page, Internet Archive capture of 22 October 2021, https://web.archive.org/web/20211022165009/https://bolt.eu/en-ng/cities/osogbo/.

Bolt Technology OÜ. 2021b. *Bolt Nigeria City Directory*. Operator city directory, Internet

Archive captures of 3 August and 26 October 2021, https://web.archive.org/web/20211026222804/https://bolt.eu/en-ng/cities/.

Church, A., M. Frost, and K. Sullivan. 2000. "Transport and Social Exclusion in London." *Transport Policy* 7 (3): 195–205. https://doi.org/10.1016/S0967-070X(00)00024-X.

Cohen, Jacob. 1960. "A Coefficient of Agreement for Nominal Scales." *Educational and Psychological Measurement* 20 (1): 37–46. https://doi.org/10.1177/001316446002000104.

Eaton, R. A., S. Gerard, and R. S. Dattilo. 1987. "A Method for Rating Unsurfaced Roads." *Transportation Research Record* (Washington, DC), no. 1106: 34–38. https://onlinepubs.trb.org/Onlinepubs/trr/1987/1106v2/1106v2-004.pdf.

Faruk, Nasir, Adeseko A. Ayeni, and Yunusa A. Adediran. 2013. "On the Study of Empirical Path Loss Models for Accurate Prediction of TV Signal for Secondary Users." *Progress In Electromagnetics Research B* 49: 155–76. https://doi.org/10.2528/PIERB13011306.

Federal Communications Commission. 2019. *Mobility Fund Phase II Coverage Maps Investigation: Staff Report*. DOC-361165A1. Federal Communications Commission, Rural Broadband Auctions Task Force and Office of Economics and Analytics. https://docs.fcc.gov/public/attachments/DOC-361165A1.pdf.

Federal Ministry of Communications and Digital Economy. 2019. *National Digital Economy Policy and Strategy (NDEPS) 2020–2030*. Federal Government of Nigeria. https://nitda.gov.ng/wp-content/uploads/2020/06/National-Digital-Economy-Policy-and-Strategy.pdf.

Federal Ministry of Communications and Digital Economy. 2020. *Nigerian National Broadband Plan 2020–2025*. Federal Government of Nigeria. https://ncc.gov.ng/sites/default/files/2024-11/Documents/Nigerian_National_Broadband_Plan_2020-2025.pdf.

GADM. n.d. *GADM Database of Global Administrative Areas, Version 4.1*. https://gadm.org.

Geurs, Karst T., and Bert van Wee. 2004. "Accessibility Evaluation of Land-Use and Transport Strategies: Review and Research Directions." *Journal of Transport Geography* 12 (2): 127–40. https://doi.org/10.1016/j.jtrangeo.2003.10.005.

Giddy, Julia Kathryn. 2019. "The Influence of e-Hailing Apps on Urban Mobilities in South Africa." *African Geographical Review* 38 (3): 227–39. https://doi.org/10.1080/19376812.2019.1589732.

Griew, Pippa, Melvyn Hillsdon, Charlie Foster, Emma Coombes, Andy Jones, and Paul Wilkinson. 2013. "Developing and Testing a Street Audit Tool Using Google Street View to Measure Environmental Supportiveness for Physical Activity." *International Journal of Behavioral Nutrition and Physical Activity* 10 (103). https://doi.org/10.1186/1479-5868-10-103.

GSMA. 2021. *The Mobile Economy Sub-Saharan Africa 2021*. GSMA Intelligence. https://www.gsma.com/mobileeconomy/wp-content/uploads/2021/09/GSMA_ME_SSA_2021_English_Web_Singles.pdf.

Haklay, Mordechai. 2010. "How Good Is Volunteered Geographical Information? A Comparative Study of OpenStreetMap and Ordnance Survey Datasets." *Environment and Planning B: Planning and Design* 37 (4): 682–703. https://doi.org/10.1068/b35097.

Herfort, Benjamin, Sven Lautenbach, João Porto de Albuquerque, Jennings Anderson, and Alexander Zipf. 2021. "The Evolution of Humanitarian Mapping Within the OpenStreetMap Community." *Scientific Reports* 11 (3037).

https://doi.org/10.1038/s41598-021-82404-z.

Houngbonon, Georges V., Erwan Le Quentrec, and Stefania Rubrichi. 2021. "Access to Electricity and Digital Inclusion: Evidence from Mobile Call Detail Records." *Humanities and Social Sciences Communications* 8 (170). https://doi.org/10.1057/s41599-021-00848-0.

International Telecommunication Union. 2021. *Measuring Digital Development: Facts and Figures 2021*. International Telecommunication Union. https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2021.pdf.

John, Jaiz. 2021. *Uber Expands to Port Harcourt, Ibadan to Provide Access to More Riders and Drivers*. *Nairametrics*, 7 October 2021, https://nairametrics.com/2021/10/07/uber-expands-to-portharcourt-ibadan-to-provide-access-to-more-riders-and-drivers/.

Kelly, Cheryl M., Jeffrey S. Wilson, Elizabeth A. Baker, Douglas K. Miller, and Mario Schootman. 2013. "Using Google Street View to Audit the Built Environment: Inter-Rater Reliability Results." *Annals of Behavioral Medicine* 45 (Supplement 1): S108–12. https://doi.org/10.1007/s12160-012-9419-9.

Kenyon, Susan, Glenn Lyons, and Jackie Rafferty. 2002. "Transport and Social Exclusion: Investigating the Possibility of Promoting Inclusion Through Virtual Mobility." *Journal of Transport Geography* 10 (3): 207–19. https://doi.org/10.1016/S0966-6923(02)00012-1.

Kuola, Omotayo Israel. 2021. *Uber Launches Ride-Hailing Operation in 2 New Cities, Ibadan and Port Harcourt*. *Technext*, 8 October 2021, https://technext24.com/2021/10/08/uber-launches-ride-hailing-operation-in-2-new-cities-ibadan-and-port-harcourt/.

Lucas, Karen, Giulio Mattioli, Ersilia Verlinghieri, and Alvaro Guzman. 2016. "Transport Poverty and Its Adverse Social Consequences." *Proceedings of the Institution of Civil Engineers – Transport* 169 (6): 353–65. https://doi.org/10.1680/jtran.15.00073.

Mahabir, Ron, Anthony Stefanidis, Arie Croitoru, Andrew Crooks, and Peggy Agouris. 2017. "Authoritative and Volunteered Geographical Information in a Developing Country: A Comparative Case Study of Road Datasets in Nairobi, Kenya." *ISPRS International Journal of Geo-Information* 6 (1): 24. https://doi.org/10.3390/ijgi6010024.

Mangla, Tarun, Esther Showalter, Vivek Adarsh, et al. 2021. *A Tale of Three Datasets: Towards Characterizing Mobile Broadband Access in the United States*. arXiv preprint arXiv:2102.07288v1. https://arxiv.org/abs/2102.07288.

Mbabazi, Emmerentian. 2019. "Impact of Unpaved Road Condition on Rural Transport Services." *Proceedings of the Institution of Civil Engineers – Municipal Engineer* 172 (4): 239–45. https://doi.org/10.1680/jmuen.18.00048.

Odusanya, Kayode, and Morakinyo Adetutu. 2020. "Exploring the Determinants of Internet Usage in Nigeria: A Micro-Spatial Approach." In *Responsible Design, Implementation and Use of Information and Communication Technology (I3E 2020)*, vol. 12066. Lecture Notes in Computer Science. Springer. https://doi.org/10.1007/978-3-030-45002-1_26.

Okoro, R. C., and P. C. Iwuji. 2019. "Investigation of GSM Signal Strength in Rural Communities in the South-Eastern Region of Nigeria." *European Scientific Journal* 15 (15): 140. https://doi.org/10.19044/esj.2019.v15n15p140.

Olowosegun, Adebola, Dumiso Moyo, and Deepak Gopinath. 2021. "Multicriteria Evaluation of the Quality of Service of Informal Public Transport: An Empirical Evidence from Ibadan, Nigeria." *Case Studies on Transport Policy* 9 (4): 1518–30. https://doi.org/10.1016/j.cstp.2021.08.002.

Ookla. n.d. *Ookla Speedtest Open Data*. https://github.com/teamookla/ookla-open-data.

OpenCellID. n.d. *OpenCellID: Open Database of Cell Tower Locations*. https://opencellid.org.

OpenStreetMap contributors. n.d. *OpenStreetMap*. https://www.openstreetmap.org.

Oughton, Edward. 2021. *Policy Options for Digital Infrastructure Strategies: A Simulation Model for Broadband Universal Service in Africa*. arXiv preprint arXiv:2102.03561. https://arxiv.org/abs/2102.03561.

Porter, Gina. 2014. "Transport Services and Their Impact on Poverty and Growth in Rural Sub-Saharan Africa: A Review of Recent Research and Future Research Needs." *Transport Reviews* 34 (1): 25–45. https://doi.org/10.1080/01441647.2013.865148.

Ravallion, Martin. 2011. "On Multidimensional Indices of Poverty." *The Journal of Economic Inequality* 9 (2): 235–48. https://doi.org/10.1007/s10888-011-9173-4.

Riddlesden, Dean, and Alex D. Singleton. 2014. "Broadband Speed Equity: A New Digital Divide?" *Applied Geography* 52: 25–33. https://doi.org/10.1016/j.apgeog.2014.04.008.

Rodríguez-Castelán, Carlos, Rogelio Granguillhome Ochoa, Samantha Lach, and Takaaki Masaki. 2021. *Mobile Internet Adoption in West Africa*. Policy Research Working Paper No. 9560. The World Bank. https://documents1.worldbank.org/curated/en/878041614611542135/pdf/Mobile-Internet-Adoption-in-West-Africa.pdf.

Salemink, Koen, Dirk Strijker, and Gary Bosworth. 2017. "Rural Development in the Digital Age: A Systematic Literature Review on Unequal ICT Availability, Adoption, and Use in Rural Areas." *Journal of Rural Studies* 54: 360–71. https://doi.org/10.1016/j.jrurstud.2015.09.001.

Sayers, Michael W., Thomas D. Gillespie, and William D. O. Paterson. 1986. *Guidelines for Conducting and Calibrating Road Roughness Measurements*. World Bank Technical Paper No. 46. The World Bank. https://documents1.worldbank.org/curated/en/851131468160775725/pdf/multi-page.pdf.

Velaga, Nagendra R., Mark Beecroft, John D. Nelson, David Corsar, and Peter Edwards. 2012. "Transport Poverty Meets the Digital Divide: Accessibility and Connectivity in Rural Communities." *Journal of Transport Geography* 21: 102–12. https://doi.org/10.1016/j.jtrangeo.2011.12.005.

World Bank. 2019a. *Nigeria Digital Economy Diagnostic Report*. No. 32743. The World Bank. https://documents1.worldbank.org/curated/en/387871574812599817/pdf/Nigeria-Digital-Economy-Diagnostic-Report.pdf.

World Bank. 2019b. *Nigeria on the Move: A Journey to Inclusive Growth. Nigeria Systematic Country Diagnostic*. The World Bank. http://documents.worldbank.org/curated/en/891271581349536392/pdf/Nigeria-on-the-Move-A-Journey-to-Inclusive-Growth-Moving-Toward-a-Middle-Class-Society.pdf.

World Bank. 2019c. *World – Measuring Rural Access: Update 2017/18*. No. ACS26526. The World Bank, Transport Global Practice. https://documents1.worldbank.org/curated/en/543621569435525309/pdf/World-Measuring-Rural-Access-Update-2017-18.pdf.

World Bank. 2020. *Nigeria Rural Access and Agricultural Marketing Project (RAAMP): Project Appraisal Document*. The World Bank. https://documents.worldbank.org/curated/en/106831580332489460.

WorldPop. 2020. *WorldPop Global 2000–2020 Constrained Individual Countries, UN-Adjusted: Nigeria 2020*. https://data.worldpop.org.