

Roodman, David

**Working Paper**

## The Arrival of Fast Internet and Employment in Africa - Comment

I4R Discussion Paper Series, No. 148

**Provided in Cooperation with:**

The Institute for Replication (I4R)

*Suggested Citation:* Roodman, David (2024) : The Arrival of Fast Internet and Employment in Africa - Comment, I4R Discussion Paper Series, No. 148, Institute for Replication (I4R), s.l.

This Version is available at:

<https://hdl.handle.net/10419/302282>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



No. 148

I4R DISCUSSION PAPER SERIES

# **The Arrival of Fast Internet and Employment in Africa Comment**

David Roodman

September 2024

## I4R DISCUSSION PAPER SERIES

I4R DP No. 148

# The Arrival of Fast Internet and Employment in Africa - Comment

**David Roodman<sup>1</sup>**

*<sup>1</sup>Open Philanthropy, San Francisco/USA*

SEPTEMBER 2024

Any opinions in this paper are those of the author(s) and not those of the Institute for Replication (I4R). Research published in this series may include views on policy, but I4R takes no institutional policy positions.

I4R Discussion Papers are research papers of the Institute for Replication which are widely circulated to promote replications and meta-scientific work in the social sciences. Provided in cooperation with EconStor, a service of the [ZBW – Leibniz Information Centre for Economics](#), and [RWI – Leibniz Institute for Economic Research](#), I4R Discussion Papers are among others listed in RePEc (see IDEAS, EconPapers). Complete list of all I4R DPs - downloadable for free at the I4R website.

I4R Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

### Editors

**Abel Brodeur**  
*University of Ottawa*

**Anna Dreber**  
*Stockholm School of Economics*

**Jörg Ankel-Peters**  
*RWI – Leibniz Institute for Economic Research*

# The Arrival of Fast Internet and Employment in Africa

## Comment

David Roodman<sup>1\*</sup>

Open Philanthropy

July 2024

### Abstract

Hjort and Poulsen (2019) frames the staggered arrival of submarine Internet cables on the shores of Africa circa 2010 as a difference-in-differences natural experiment. The paper finds positive impacts of broadband on individual- and firm-level employment and nighttime light emissions. These results largely are not robust to alternative geocoding of survey locations, to correcting for a satellite changeover at end-2009, and to revisiting a definition of the treated zone that has no clear technological basis, is narrower than the spatial resolution of nearly all the data sources, and is empirically suboptimal as a representation of the geography of broadband.

### Introduction

Hjort and Poulsen (2019)—HP19 for short—is the first work to estimate the distributional impacts in poor countries of a technology thought to have increased inequality in rich ones (HP19, p. 1036). HP19 obtains its evidence by creatively

\* david.roodman@openphilanthropy.org. Thanks to Jonas Hjort for assistance with data and code and feedback on an earlier draft, to Otis Reid for guidance and ideas, and Steve Song for advice. The packages `esttab` (Jann 2007), `blindschemes` (Bischof 2017), `grc1leg2` (Over 2022), `reghdfe` (Correia 2014), and `reghdfejl` (Roodman 2024) were used to estimate, tabulate, and plot. The nighttime light data used here is produced by Earth Observation Group, Payne Institute for Public Policy, Colorado School of Mines, using DMSP data collected by US Air Force Weather Agency. 3G coverage data © Collins Bartholomew and GSMA, used with permission. Replication data and code at <https://doi.org/10.7910/DVN/V7ASIP>.

framing a difference-in-differences (DID) design around the landing of ten under-sea data cables on the shores of Africa in 2009–12. Before then, the continent’s Internet bandwidth was constrained by older cables and satellite links. The removal of the bottlenecks surely delivered higher data speeds to households and firms served by national fiber optic networks. The cable connection events thus generated treatment variation over both space and time. Given the uncertainties in large engineering projects, the timing component in particular was plausibly, *locally* exogenous—exogenous, that is, at the scale of a year.

HP19 traces the impacts of this treatment in many data sets, nearly all of which are gathered here too: city-level bandwidth measurements from the content delivery company Akamai; multi-country household surveys from Afrobarometer and the Demographic and Health Surveys (DHS) program; South Africa’s Quarterly Labour Force Surveys (QLFS); Ethiopia’s firm-level Large and Medium Scale Manufacturing Industries Survey (LMMIS); the World Bank’s multi-country enterprise surveys (WBES); and nighttime light observations from the Defense Meteorological Satellite Program’s Operational Linescan System (DMSP).

HP19 finds positive impacts from broadband access in all the data sets. The collective coherence of the results adds to their individual credibility. Local broadband service increases the probability of an individual being employed by 4.6 percentage points in the DHS follow-ups, 7.7 points in the Afrobarometer surveys, and 2.2 in the South Africa QLFS. The effect is positively skill-biased, meaning that it is largely confined to employment in skilled occupations. Yet, counterintuitively, the effect size is comparable for less- and more-schooled individuals. Separately, broadband access is associated with 2–3% higher nighttime light emissions, which is taken to indicate increased economic activity (Chen and Nordhaus 2011; Henderson, Storeygard, and Weil 2012). HP19 also finds positive impacts on firm entry and firm exports.

Much of the work in HP19 relates to the geographic dimension of the data:

constructing maps of backbone networks circa 2010, estimating coordinates for localities identified in survey data only by place name, and classifying locales as “treated” or “untreated” according to proximity to digital infrastructure.

This comment mainly revisits that geographic work. One step I take is to re-estimate the geographic coordinates of observations in primary sources whose locations are given by name. I recode the Akamai and Ethiopia LMMIS data using public sources such as Google Maps. For Afrobarometer enumeration areas, I turn to the third-party geocoding by AidData (BenYishay et al. 2017).

I also revisit the geographic definitions of the treatment and control groups. HP19’s base specifications define the treatment zone as the area within  $0.005^\circ$  of a fiber optic line, which is about 560 meters. (HP19 does vary this parameter in a robustness test.) The control zone is the area within  $0.005$ – $0.1^\circ$ , or up to about 11 kilometers. The  $0.005^\circ$  threshold invites critique from a few directions. It is not rooted in technical characteristics of communications equipment. Both wired and wireless means deliver broadband more than 560 meters from the backbone in many cities. And the threshold produces an asymmetry, a control zone 19 times wider than the treatment zone. The precision in the definition is in a sense illusory since almost none of the location measurements in the HP19 data, including the paths of the fiber lines are reliably estimated to the sub-kilometer scale. As mentioned, many of the primary sources identify locations only by naming localities or districts, which can be well more than 560 meters wide. The DHS data includes numerical latitude and longitude coordinates, but injects kilometer-scale noise for privacy. Even the satellite data may contain comparable measurement error (Elvidge et al. 2004), and at any rate HP19 downsamples the light data to  $0.1^\circ$  squares. In effect, by preferring a narrow treatment band, HP19 attains its results after reducing the power to detect them. To the extent that the research and public process favor statistically significant results—however unconsciously—lower power is associated with higher propensity to generate significant but non-representative results (Ioannidis, Stanley, and

Doucouliaos 2016). And while the measurement error attenuates exogenous identifying variation, creating a conservative bias, it does not attenuate endogenous influences in tandem, so they gain relative importance.

Acting on suggestions from the experts who supplied the underlying data on the geography of the African backbone network, several alternative representations are proposed here, and tested for fidelity: defining the treatment zone as  $0.05^\circ$  from fiber lines instead of  $0.005^\circ$ , while retaining the  $0.1^\circ$  radius for the control zone bound; doubling both those numbers, to  $0.1^\circ$  and  $0.2^\circ$ ; and retaining the latter bounds while taking distance from the nearest network *node* rather than the nearest line.<sup>2</sup> The motivation for taking distance from nodes is that HP19's line-based representation causes localities in "digital flyover country" to be coded as treated because fiber lines run nearby, regardless of whether they provide broadband locally.

Empirically, the alternatives more accurately represent the geography of broadband. For example, using an early (2014) mobile phone 3G coverage map, the  $t$  statistic for a difference in broadband access across the HP19-preferred treatment and controls zones is 1.82 in the DHS sample. For the HP19-geocoded Afrobarometer locations, the  $t$  is  $-4.00$ : the rate of 3G coverage was *lower* in the HP19 treated sample. Switching to the  $0.1^\circ$  and  $0.2^\circ$  bounds for treatment and control zones (and to the Afrobarometer-endorsed geocoding) produces corresponding  $t$  statistics of 11.03 for DHS and 9.09 for Afrobarometer. That indicates that the wider zones represent the geography of broadband with more fidelity, and so should undergird higher-powered impact estimation.

Lastly, this reanalysis introduces two changes in the DMSP-OLS nighttime

<sup>2</sup> In conversations, Paul Hamilton of Africa Bandwidth Maps, and Steve Song, creator of the AfTerFibre fiber network map, preferred defining treatment on the basis of distance from network nodes rather than fiber lines, and bounding the treatment zone at  $0.05^\circ$  or  $0.1^\circ$  rather than  $0.005^\circ$ . In light of the diversity of local circumstances and incompleteness of the available data, neither endorsed these proposals with high confidence.

light data: taking it at its original  $1/120^\circ$  resolution instead of a down-sampling it to  $0.1^\circ$ ; and switching to an “intercalibrated” data set that partially removes time effects such as large discontinuities from a satellite changeover between 2009 and 2010.

Most HP19 findings shrink in magnitude and statistical significance under these revisions. A meta-analytic average of the three estimates of the impact on individual-level employment (using DHS, Afrobarometer, and South Africa QLFS data) drops from 3.6 points (standard error 1.2 points) for the original results to 2.6 (0.7) after making a few data and specification corrections and revising the geocoding in some data sets. The average employment impact falls to 0.4 (0.7) or 0.5 (0.7) points when using higher-fidelity treatment and control zone classifications. The main source of positivity in these results is the DHS data. Yet that family of surveys is the least frequent of the ones used here, with an average gap between pre- and post-event surveys of almost 9 years. Its identifying temporal variation is thus dominated by long-term trends that are, while more important than short-term changes, also more presumptively endogenous.

As to firm-level impacts, in the regressions on World Bank Enterprise Survey data, identification turns out to flow solely from the Nigeria data, which are geographically labeled only by state and thus contain almost no information about proximity to fiber networks. The Ethiopia firm-level results are robust to re-geocoding, but not to widening the treatment and control zones.

In the nighttime light data, the impact of broadband access on emissions appears mainly to be an artifact of a satellite changeover between 2009 and 2010.

Section 1 reviews the HP19 specifications. Section 2 discussed the HP19 definitions of treatment and control zones and details the re-geocoding of three data sets. Sections 3 and 4 present results for non-satellite and satellite data. Section 5 concludes.



## 1 The HP19 specifications

HP19 estimates impacts on a variety of outcomes through specifications of a broadly shared form, linear two-way fixed effects (TWFE). The predictor of interest is the product of temporal and cross-sectional indicators, namely, a dummy for whether an observation occurred after a country received undersea cable connection between 2009 and 2012 and a dummy for whether the observation occurred proximate to a fiber line thought active circa 2010. The idea is that, within each country, broadband access was distinctly greater after a connection event, more so for locales close to the national fiber backbone.

The data sets to which TWFE is applied differ in their cadence. The multi-country DHS, Afrobarometer, and WBES surveys are infrequent in any given country. So from these sources HP19 uses one “before” and one “after” round for each country.<sup>3</sup> The average of the country-level before-after gaps is almost 9 years for the DHS, 4 for Afrobarometer, and immaterial for the WBES, as will be explained. The Ethiopia firm survey and night light data are annual, and the Akamai Internet speed and South Africa labor force data quarterly. From these higher-frequency sources, multiple rounds are taken.

The HP19 specifications control for several sets of fixed effects. Since the predictor of interest is the product of geographic and temporal indicators, the TWFE design requires the control set to cover those two factors separately, with sets of fixed effects. When working with the pooled data from multi-country surveys, HP19 goes farther, interacting the temporal indicators (year, quarter, or post-treatment dummies) with country dummies. The regressions thus pool country-specific TWFE quasi-experiments while subjecting them to the restriction of equal impacts. To reduce spatial autocorrelation, HP19 typically adds more fixed effects, for enumeration areas (South Africa QLFS and Ethiopia LMMIS), cities (Akamai and WBES), or interactions between the treatment

<sup>3</sup> There is one exception: the WBES data includes three rounds from Mali.

dummy and  $0.1^\circ \times 0.1^\circ$  grid cell dummies (DHS and Afrobarometer).

Pursuant to its Data Availability Policy, the *American Economic Review* received and posted data and code needed to generate the HP19 results. Also pursuant to the policy, the HP19 authors have provided intermediate data and code, in response to a request made in 2022. However, there are gaps. For example, my geocodings of the Akamai and Ethiopia LMMIS data differ from the original in some cases, for reasons that cannot be fully determined from the public data and code. Details of the construction of the HP19 fiber backbone map have also not been shared.

## 2 Revisiting the geographic information and definitions

### 2.1 Reconstructing the distance variable

I obtain DHS, Afrobarometer, WBES, DMSP, and South Africa QLFS data sets from their primary providers. I take Akamai and Ethiopia LMMIS data from the HP19 archive. I revisit the geographic processing of all of the data sets. Several revisions result:

1. The AidData geocoding of the Afrobarometer household survey locations (BenYishay et al. 2017) are used in the place of the HP19 geocodings. AidData, based at the College of William and Mary, has long experience in geocoding data from developing countries. The coding project was carried out by a “team of trained geocoders” according to a “double-blinded” process in which two people independently coded each location, disagreements went to an arbitration round, and quality was checked in other ways. Here, only geocodes scored by BenYishay et al. as fully precise (precision level 1) are retained.
2. The 77 localities sampled by the Ethiopia LMMIS are manually re-geocoded by drawing on publicly available GIS data on towns and woredas

(third-level administrative units).<sup>4</sup> The LMMIS data set in the HP19 public archive names all three administrative levels for each observation.

3. Most localities contributing identifying information in the Akamai Internet speed data set are similarly recoded. These localities are those that satisfy HP19's criteria of providing more than 10 speed measurements in at least one quarter; and being in countries with data adequate for TWFE, i.e., ones with at least one treatment city and at least one control city with pre- and post-treatment data. OpenStreetMap's Nominatim engine is the primary source of coordinates, and Google Maps a secondary one.<sup>5</sup> 167 of the 186 localities are coded. This coding is less complete and confident than some others because only country and city names are available in the HP19 public archive. (Precisely this issue, in the WBES data, leads HP19 to deemphasize results from that data set.) A few Akamai city names in the HP19 archive are ambiguous, perhaps because of truncation; one in Benin is "De," for example. Other ambiguities arise for lack of state names: at least four villages in Nigeria are named "Asa."<sup>6</sup> A few other "cities" are states or provinces and are left uncoded because of their geographic coarseness.<sup>7</sup>

<sup>4</sup> Town coordinates are taken from the Google Maps API. For Woredas, coordinates of prominent towns on major roads are used, since they are better guesses for the locations of large and medium-sized firms.

<sup>5</sup> The Nominatim engine often returns multiple results. In most cases, I use the result having *class* = "place" with highest *importance*. In the case of ties—typically Nigerian villages with the same name—I do not geocode unless Google Maps favors one locality.

<sup>6</sup> See [nominatim.openstreetmap.org/search?q=asa&countrycodes=NG](https://nominatim.openstreetmap.org/search?q=asa&countrycodes=NG).

<sup>7</sup> I also revisit the list of top-four largest cities in each country, which HP19 exclude from some Akamai regressions because of ambiguity in their treatment classification. I add Porto-Novo, Benin; Kisangani, DRC; Tema, Ghana; Makuru, Kenya; Nampula, Mozambique; and Dodoma and Zanzibar, Tanzania. I drop Arusha, Tanzania; and Bloemfontein and East London, South Africa. The source for all changes

4. The DMSP light data is taken at its native resolution of  $1/120^\circ$  rather than downsampled to  $1/10^\circ$ . An intertemporally calibrated version of the data is also analyzed, as explained in section 4.

## 2.2 Defining treatment and control groups

As in the human body, a fiber optic backbone supports a filigreed network. At junctions called nodes, feeders branch out from the backbone, in a variety of technological forms. These technologies include asymmetric digital subscriber line (ADSL), which runs over copper phone lines; Data Over Cable Service Interface Specification (DOCSIS), which does the same on cable television wires; local fiber lines (“fiber to the home”); fixed wireless, which is essentially high-powered Wi-Fi; WIMAX, the mesh-network equivalent; microwave, for long-distance transmission; and cellular towers, which are considered to supply broadband in generation 2.5G or higher and which are themselves linked to the backbone through wired or wireless means. Most broadband in Africa is delivered wirelessly. In 2012, Africa had 2 million wired broadband subscriptions and 51 million wireless ones. By 2021, the totals had risen to 7 million and 448 million.<sup>8</sup>

### The HP19 treatment and control zone definitions

To estimate the impacts of broadband in Africa under a dichotomous definition of connectivity as in HP19, one needs to classify observations according to whether they have access to high bandwidth in the post-treatment period. It is not a straightforward question how best to do so, given the complex and rapidly changing anatomy of the continent’s networks, and the limited data thereon circa 2010. This section describes HP19’s approach, comments on it, and

is, as in HP19, UNdata ([data.un.org/Data.aspx?d=POP&f=tableCode%3a240](http://data.un.org/Data.aspx?d=POP&f=tableCode%3a240), using pre-2015 observations) except for the DRC, which is absent from that source, and for which GeoNames is used ([geonames.org/CD/largest-cities-in-.html](http://geonames.org/CD/largest-cities-in-.html)).

<sup>8</sup> See International Communications Union data set, [itu.int/en/ITU-D/Statistics/Documents/facts/ITU\\_regional\\_global\\_Key\\_ICT\\_indicator\\_aggregates\\_rev1\\_Jan\\_2022.xlsx](http://itu.int/en/ITU-D/Statistics/Documents/facts/ITU_regional_global_Key_ICT_indicator_aggregates_rev1_Jan_2022.xlsx), “By BDT region” tab.

proposes alternatives. It will be argued that the alternatives more accurately represent the geography of broadband in the study setting, and so endow impact regressions with greater power. And switching to evidence-backed alternatives tests robustness in a minimally arbitrary way.

HP19's classification proceeds in three steps. First, the decision is made to proxy broadband availability by *distance from fiber infrastructure*. Finer-grained information on mobile phone coverage or distance from municipal fixed-line networks would be preferable, but is scarce for the study setting.<sup>9</sup> (For many African countries, 3G coverage data only begins in 2014.) Second, fiber optic *lines* are chosen as the distance referent. Third, *bounds* (radii) are set to demarcate the treatment and control groups. The treatment zone is set to 0.005 angular degrees from the lines and the control zone from there to 0.1°. At the Earth's surface, each angular degree extends about 111.32km, so 0.005° and 0.1° are approximately 0.56km and 11km.

One downside of HP19's second step, taking fiber lines as the referent, is that some people live and work in digital flyover country—near the lines, but not connected to them. Metaphorically speaking, they can watch the bits fly by, but not receive or send them. *Access* requires connection to a network node via wired or wireless feeders. Since Africans usually obtain broadband through cell towers, many of which are erected near fiber nodes, distance to the nearest node may correlate better with bandwidth availability. For this reason, I introduce a geo-coded fiber node data set from Africa Bandwidth Maps (ABM).

Whether a given node-based representation coheres better with the

<sup>9</sup> Related studies, in later years or other countries, use cellular coverage data. Collins Bartholomew coverage maps feature in Manacorda and Tesei (2020), Masaki, Ochoa, and Rodríguez-Castelán (2020), Gonzalez (2021), Guriev, Melnikov, and Zhuravskaya (2021), Chiplunkar and Goldberg (2022), Goldbeck and Lindlacher (2022), and Mensah, Tafere, and Abay (2022). Klonner and Nolen (2010) uses a South Africa-specific source. Bahia et al. (2020) models coverage using information on cell towers.

geography of broadband than a given line-based representation is an empirical question, which will be explored just below. Many cell towers, which are the true points sources for most broadband in Africa, are not right at fiber nodes. Meanwhile, if households and firms congregate along the major roads that fiber lines also follow, and their purchasing power attracts cell towers, a line-based representation of broadband access may in fact be reasonably accurate.

With more confidence, this reanalysis proposes an alternative to the choices made in the third step of the HP19 classification, the delimitation of the treatment and control zones by  $0.005^\circ$  and  $0.1^\circ$  radii. The initial version of HP19 places no outer bound on the control zone and sets the treatment boundary at the *median* distance from the network, which is about 19.7km for the Akamai data set, 37.5km for DHS, 2.7km for Afrobarometer, 1.04km for the Ethiopia LMMIS, 3.4km for the South Africa QLFS, and 3.8km for WBES (Hjort and Poulsen 2016, note 36). The final version (HP19, note 34) justifies the switch to the much narrower  $0.005^\circ$  treatment bound as follows: “Technological considerations indicate that 500 meters is a reasonable proxy for potential fast Internet reach beyond the backbone cables for copper-cable last mile technologies.”

The  $0.005^\circ$  marker appears hard to justify. HP19 does not specify the technological considerations it mentions. Perhaps the reasoning behind the emphasis on copper in HP19’s explanation is that wired connections, where available, exercise disproportionate economic influence. But, as already noted, wireless access dominates in Africa. It appears central to three of HP19’s five anecdotal examples of impact (HP19 §I.C): rising eCommerce, domestic manufacture of low-end Internet-capable devices, and improved coordination among farmers and other supply chain actors. To the extent that the reach of copper networks matters, copper can in fact reach much farther than 500 meters, as it does in cities with wired phone networks. And even if it could not, when the outcome is individual-level employment and the place of observation is the home, the economic influence of broadband would spread farther: broadband would reach firms within

500 meters, which would hire people living more distal from the lines.

A final concern is that the use of a  $0.005^\circ$  ( $\sim 0.56\text{km}$ ) bound implies an unrealistic expectation of precision in the data. Each download speed, household, or firm observation in HP19 is geocoded by an enumeration area or some larger locality such as a city, to which a representative coordinate pair is assigned. Most of the localities themselves extend more than  $0.56\text{km}$  in one direction or another, the main exception being certain urban enumeration areas in the South Africa QLFS that are as small as a building. Since all observations in each enumeration area or city are assigned the same coordinate, all are assigned the same treatment status, even when some might be considered treated and some not, if their true distances from fiber infrastructure were known. Distinctively, the DHS provides actual coordinates for the centroids of enumeration areas, not just place names. But it randomly perturbs the coordinates for privacy—by  $0\text{--}2\text{ km}$  for urban observations,  $0\text{--}5\text{km}$  for 99% of rural ones, and  $0\text{--}10\text{km}$  for the other 1% (Perez-Heydrich et al. 2013).<sup>10</sup> Even the satellite data may contain kilometer-scale error: Elvidge et al. (2004, p. 290) finds positional errors of  $1.55\text{--}2.36\text{km}$  in DMSP readings, at least when taken from a single orbit, rather than averaging together many as is done for the annualized DMSP data used in HP19. At any rate, HP19 downsamples the light data to  $0.1^\circ$  ( $\sim 11\text{km}$ ) resolution, from its native resolution of  $1/120^\circ$ . Finally, the fiber paths are digitized with error too. Since more than 90% of fiber lines in Africa run along roads,<sup>11</sup> the routes in the AfTerFibre database that is the basis for the HP19 network map are usually derived from vector representations of major roads, which do not capture all bends

<sup>10</sup> The DHS perturbs reported survey locations by drawing from uniform distributions for the angle and distance of perturbation. Under this process, if the maximum perturbation distance is  $R$  then the average absolute perturbation along a given axis is  $R/\pi$ . If a survey location were exactly on a fiber line, if that line's path were locally straight, and if the recorded path did not also contain measurement error, then the expected estimated distance from point to line would be  $R/\pi$ .

<sup>11</sup> Conversation with Paul Hamilton, Africa Bandwidth Maps, August 30, 2022.

and curves.<sup>12</sup> Figure A–1 in the appendix illustrates with the case of a line running through Cotonou, Benin. As digitized, the line deviates up to a kilometer from the curved national highway that in reality it probably follows.

Giving the treatment zone a width at or below the resolution of most of the geographic data injects substantial error into the HP19 treatment assignments. One might reasonably share HP19’s expectation (note 34) that this error makes the estimation strategy conservative. Measurement error in a right-side variable attenuates its coefficient estimate. However, that result assumes that treatment and the measurement error are exogenous. In observational data on complex social systems, in which causal arrows run in all directions, such assumptions are not strictly true, even after double-differencing and adding controls. In general, the pragmatic hope is that in a well-framed design, biases from violations of identifying assumptions are second-order. However, as the treatment zone narrows beyond the resolution of the data, erstwhile second-order biases can gain relative importance. Appendix B demonstrates this claim rigorously. Moreover, reducing power can combine in a distortionary way with conscious or unconscious biases within the research and publication process toward statistically significant results, raising the risk that published results are unrepresentative of the sampling distributions they are described as coming from, and thus potentially fragile (Ioannidis, Stanley, and Doucouliagos 2016).

### **Alternative treatment and control zone definitions**

To ground an exploration of alternative bounds, I first illustrate the geography of broadband in Africa with two maps, one of a city, one of a country. Figure 1 depicts the Nairobi region. A burst of technology start-ups there constitutes one success story in HP19, in which wired broadband was plausibly crucial. The top layer of the map, in orange squares, shows where the Internet speed testing

<sup>12</sup> AfTerFibre database, [AfTerFibre.carto.com/tables/af\\_fibrephase/public](https://af-terfibre.carto.com/tables/af_fibrephase/public), accessed June 11, 2022.



company Ookla measured wired connections exceeding 256 kilobits per second (Kbps), a common definition of broadband. The data source starts in the first quarter of 2019, making that the period closest to the HP19 study setting, within the limits of data availability.<sup>13</sup> A layer down, in green, are locations of Ookla *wireless* broadband measurements. The lowest layer shows 3G coverage from a different data set, which is maintained by Collins Bartholomew on behalf of the international industry association GSMA. Mobile phone network operators supply data to it voluntarily—or don't. Substantial information first appears in HP19 study countries, aside from South Africa, in 2014, so that year's values are mapped. 2014 is also the follow-up year for Kenya in HP19's DHS sample. Collins Bartholomew asks network operators to distinguish "strong" coverage (signal strength above -92 decibels) from "variable" coverage, which may be unreliable indoors. Those coverage levels are shown in dark and light purple. The distribution of strong 3G coverage is taken to proxy availability of wireless broadband, as 2G is much slower and 4G was hardly available circa 2010 outside of urban South Africa.

The two data sources on the reach of broadband may be biased in opposite directions. The Ookla data set is a convenience sample, and so may understate access to broadband. Mobile network operators have an incentive to exaggerate the reach of their towers.

Finally, the map depicts the routes of fiber optic lines as recorded in the public HP19 data archive. The thin yellow ribbons with dashed borders show the HP19 0.005°-radius treatment zone while the wider yellow ribbons show the 0.1°-radius control zone.

The map of the Nairobi region supports certain conclusions. First, technological limits on wired connections did not confine broadband to within 0.005° of fiber lines. Indeed, HP19's control zone bound of 0.1° embraces much territory

<sup>13</sup> Speedtest by Ookla Global Fixed and Mobile Network Performance Maps, [registry.opendata.aws/speedtest-global-performance](https://registry.opendata.aws/speedtest-global-performance), accessed May 4, 2023.

apparently reached by broadband in 2014. Second, the area of intensive broadband availability centers not so much on fiber lines as on a major junction in the network. To more accurately represent the distribution of broadband access with dichotomized treatment and control zones, it would seem reasonable to widen the bounds from  $0.005^\circ$  and  $0.1^\circ$  to  $0.1^\circ$  and  $0.2^\circ$ , and to switch their referent from network lines to network nodes. These wider bounds would have the added virtues of not implying an unrealistic demand for precision in the available data, and being minimally arbitrary, round numbers that give the same width to the treatment and control zones.

The second map, of Ghana, provides a larger-scale check on the impressions from the first. It also adds blue circles to represent the  $0.1^\circ/0.2^\circ$ -from-nodes classification just proposed, using the list of nodes from Africa Bandwidth Maps. (See Figure 2.) Again, HP19's treatment zone—depicted as the thin yellow spines within the wider yellow bands—appears to exclude much treated area. Yet the node-based definition also commits many type I and type II errors. If anything, a point-source-based representation would be more accurate *away* from fiber lines (and fiber nodes), where service is provided by isolated cell towers uplinked by microwave or satellite—assuming the data on point sources included such locations. Along the southern coast of Ghana, the range of 3G coverage parallels the fiber lines and is not confined to areas near nodes recorded in the data used here.

To check these casual impressions, I compute, for each of several possible definitions of the treatment and control zones, and each broadband proxy (Ookla measurements and “strong 3G” coverage from the GSMA), the  $t$  statistic for the hypothesis that there is no treatment-control difference in the rate of broadband access. I compute the statistic separately for the geocoded locations in all four data sets in which employment effects will be estimated. Samples are subject to HP19 restrictions, such as age being below 65. The “strong 3G” calculations include only countries with substantial 3G data as of 2014, which are

Ghana, Kenya, Madagascar, South Africa, and Tanzania. The  $t$  statistics are computed for five classifications: the original HP19 data and definitions (0.005/0.1° from lines); a revised version that introduces certain geocoding and sample changes while retaining the HP19 radii (see section 3.3); and three alternatives, namely, 0.05°/0.1° from lines, 0.1°/0.2° from lines, and 0.1°/0.2° from nodes.<sup>14</sup> The last only factors in nodes active in 2010.<sup>15</sup>

The numerical results, all in Table 1, show that the alternatives improve on the HP19 classification. In the HP19 Afrobarometer sample, under the HP19 geo-codings, the  $t$  statistics are negative, meaning that the control zone had more broadband (last column of Table 1). Re-geocoding Afrobarometer place names improves matters, to the extent of bringing the  $t$  statistics close to zero (second row of the table). The 0.05°/0.1°-from-lines definition performs better in all data sets except the Ethiopia LMMIS. Doubling to the 0.1°/0.2°-from-lines classification produces the largest  $t$  statistics in all data sets save the Ethiopia one. The 0.1°/0.2°-from-nodes classification performs not quite as well, the exception again being Ethiopia. In the rest of this paper, I will for concision omit impact estimates using 0.05°/0.1°-from-lines. I will prefer results from the two final (and overall best-performing) classifications, 0.1°/0.2° from lines and nodes.

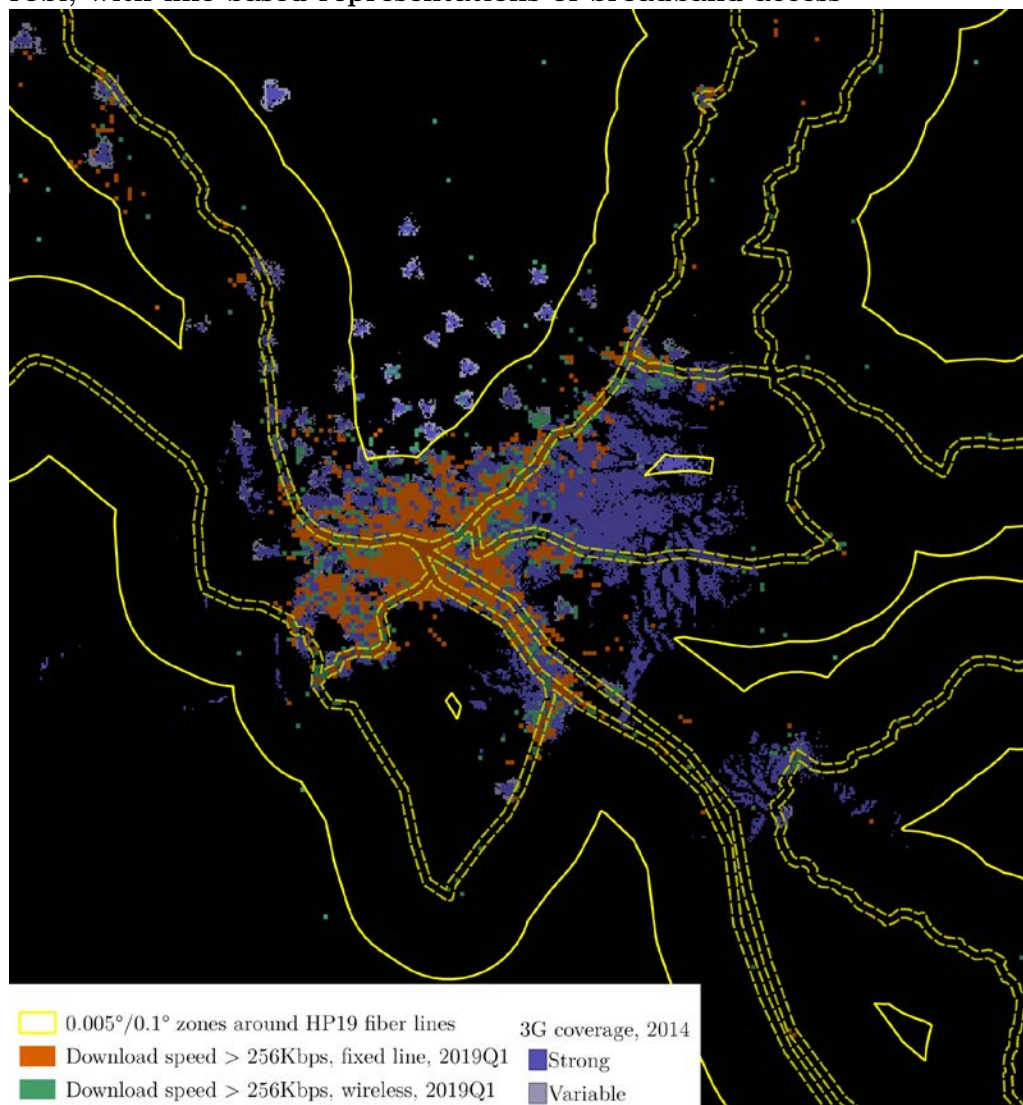
Two objections might be raised against this preference, though neither with much force. First, while a treatment classification that represents actual broadband access with more fidelity should increase power, the added power might come at a price in bias, as is often the case in econometrics. But there is little reason to expect that comparisons between the 0–0.1° and 0.1–0.2° zones should contain more endogeneity bias than ones between the 0–0.005° and 0.005–0.1°

<sup>14</sup> The  $t$  statistics are clustered in the same way as in corresponding regressions in section 3.4.

<sup>15</sup> HP19 models the fiber network as of 2009, which argues for using the node list of that year as well. However, 2010 is the first year in which all study countries—notably the DRC—have nodes in the ABM data.

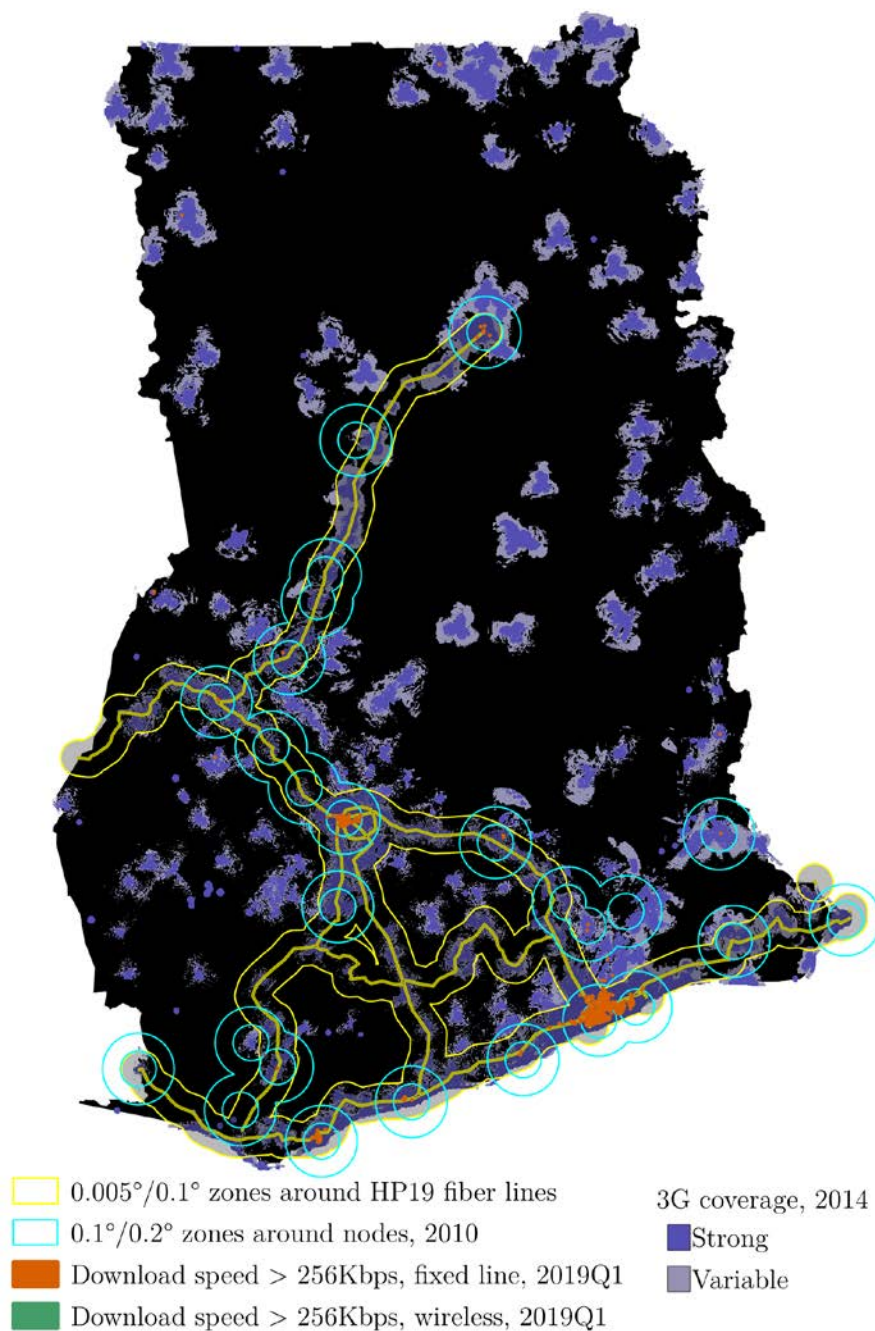
zones. Similarly, it might be objected that while the new classifications have higher fidelity in the cross-section, the same does not necessarily hold in the TWFE design. Perhaps the bandwidth *added* circa 2010 still accrued mainly within  $0.005^\circ$  of the fiber lines. But the coverage data benchmarked against here pertain to years just after treatment arrived, and the maps make it appear unlikely that technological factors limited the reach of additional broadband to  $0.005^\circ$  from fiber lines.

**Figure 1. 3G coverage and points of high download speed around Nairobi, with line-based representations of broadband access**



Notes: The base map shows, in shades of purple, 3G coverage in 2014. Orange (green) dots show locations with fixed (wireless) download speeds above 256 kilobits in the first quarter of 2019, according to Ookla. Yellow lines depict the HP19 representation of connectedness.

**Figure 2. 3G coverage and points of high download speed in Ghana, with line- and node-based representations of broadband access**



Notes: Symbology is the same as in Figure 1, with the addition of blue circles to present node-based representation of broadband access. Node locations, for 2010, are from Africa Bandwidth Maps.

**Table 1. Coverage rates of broadband under various proxies for broadband and definitions of treatment and control groups**

Specification	Broad-band proxy	DHS					Afrobarometer				
		Treatment		Control		<i>t</i>	Treatment		Control		<i>t</i>
		Rate	<i>N</i>	Rate	<i>N</i>		Rate	<i>N</i>	Rate	<i>N</i>	
Original, 0.005°/ 0.1° from lines	GSMA	0.67	3,419	0.57	17,210	1.82	0.52	1,257	0.69	3,007	-4.00
	Ookla	0.25	5,658	0.14	30,917	3.12	0.22	1,622	0.28	4,108	-1.89
+ Re-geocode, recalculate distance	GSMA	0.47	7,651	0.31	40,383	4.34	0.75	1,021	0.76	3,126	-0.23
	Ookla	0.18	12,687	0.10	67,762	4.07	0.29	1,405	0.27	4,262	0.51
+ Change to 0.05°/ 0.1° from lines	GSMA	0.40	36,582	0.12	11,452	10.90	0.76	3,539	0.72	608	0.87
	Ookla	0.14	62,305	0.02	18,144	12.13	0.29	4,888	0.15	779	4.56
+ Double to 0.1°/0.2° from lines	GSMA	0.33	48,034	0.10	17,568	11.03	0.75	4,147	0.33	661	9.09
	Ookla	0.11	80,449	0.00	25,702	14.09	0.27	5,667	0.08	1,028	7.77
+ Change to 0.1°/0.2° from node	GSMA	0.50	27,152	0.20	20,877	10.70	0.86	2,839	0.64	1,066	5.22
	Ookla	0.15	52,211	0.04	30,695	9.79	0.33	4,028	0.15	1,622	6.51
		South Africa QLFS					Ethiopia LMMIS				
Original, 0.005°/ 0.1° from lines	GSMA	0.97	24,168	0.91	155,477	5.89					
	Ookla	0.31	24,168	0.21	155,477	3.42	0.82	2,496	0.01	453	5.51
+ Re-geocode, recalculate distance	GSMA	0.97	17,514	0.91	155,993	4.63					
	Ookla	0.30	17,514	0.21	155,993	2.80	0.76	2,758	0.10	412	3.42
+ Change to 0.05°/ 0.1° from lines	GSMA	0.93	124,718	0.87	48,789	3.75					
	Ookla	0.24	124,718	0.16	48,789	4.24	0.68	3,139	0.00	31	3.29
+ Double to 0.1°/0.2° from lines	GSMA	0.92	173,507	0.71	28,045	7.75					
	Ookla	0.22	173,507	0.11	28,045	6.47	0.68	3,170	0.00	4	3.23
+ Change to 0.1°/0.2° from node	GSMA	0.93	116,898	0.83	64,482	6.77					
	Ookla	0.26	116,898	0.16	64,482	5.63	0.77	2,738	0.00	178	4.28

Notes: The table shows the fraction of various data sets (not) covered by broadband, restricting to observations that contribute identification in later regressions. Two proxies for broadband are used: “strong” 3G coverage as of 2014, according to data from the GSMA, which covers Ghana, Kenya, Madagascar, Nigeria, South Africa, and Tanzania; and Ookla-measured download speeds in the first quarter of 2019. *t* statistics are for the hypothesis of no difference in coverage between treatment and control groups, and indicate the fidelity of the HP19 treatment/control grouping to actual broadband access. In the first pair of rows in each subtable, variables and samples are defined as in HP19; in particular, outer treatment and control radii are 0.005° and 0.1°. Subsequent pairs accumulate changes. In the second pair, variables are revised as described in the notes to Table 3 and Table 4. In the third, the treatment radius is increased to 0.05° as in HP19 (§IV.C). In the fourth, radii are doubled to 0.1° and 0.2°. In the last, network nodes are the distance referent. Standard errors clustered by locality are in parenthesis.

### 3 Results from non-satellite data

#### 3.1 Geocoding

Recall that the Akamai, Ethiopia LMMIS, and Afrobarometer data sets are here

recoded from place names in the primary sources—the last through recourse to the AidData effort (BenYishay et al. 2017). It is also for Afrobarometer that I have the HP19 longitude and latitude estimates. Figure A–2 in the appendix therefore compares the HP19 and AidData coordinate assignments by drawing lines between corresponding coordinate pairs and placing dots on the AidData ends. Seven of the nine countries have at least one locale that HP19 places outside the nation’s borders. These evident errors may not affect estimates much since they are few. But they suggest that AidData checked quality more thoroughly.

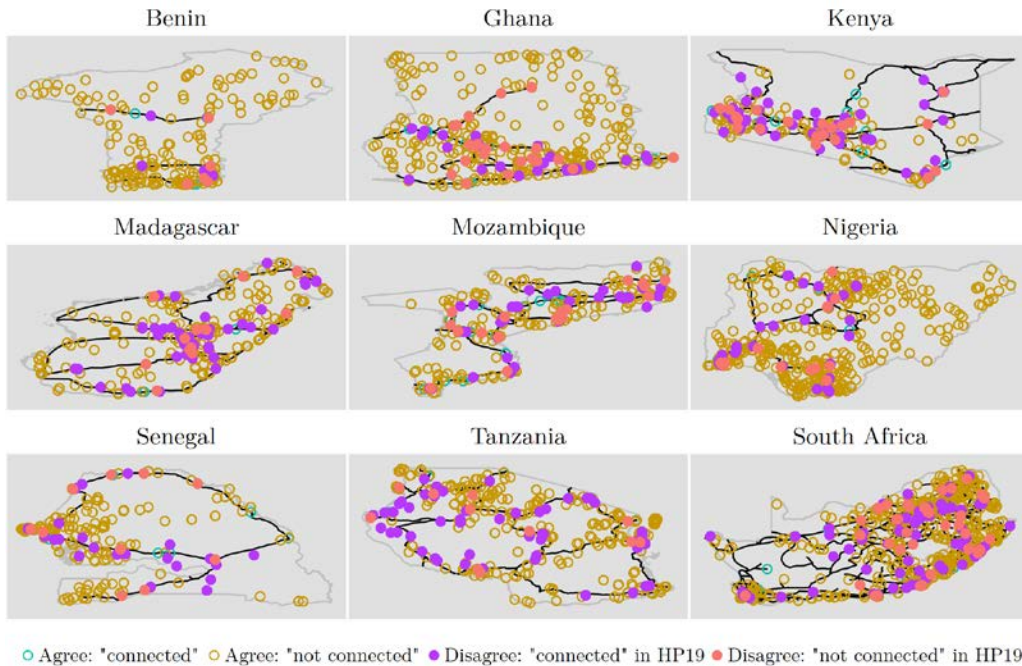
To assess the impact of the geocoding revisions on treatment classification of the Afrobarometer observations, Figure 3 plots each survey locale at its AidData location, using symbols to indicate agreement or disagreement between old and new codings under HP19’s preferred  $0.005^\circ/0.1^\circ$ -from-lines classification. Figure A–3 through Figure A–6 in the appendix do the same for the Akamai, DHS, South Africa QLFS, and Ethiopia LMMIS data sets.<sup>16</sup> Discrepancies are non-existent for the DHS and unusual for the QLFS, presumably because these geocodings allow the least researcher discretion: the DHS locations come directly from the data provider while the QLFS locations are centroids of polygons provided by Statistics South Africa. The HP19 and new treatment-control classifications agree on 68.6% of locales for Akamai, 100% for DHS, 70.4% for Afrobarometer, 93.7% for the South Africa QLFS, and 83.5% for the Ethiopia LMMIS.<sup>17</sup>

<sup>16</sup> The HP19 versions of these data sets that are available to me do not include geographic coordinates, but do include a distance variable sufficient for assigning treatment status.

<sup>17</sup> “Classification” here is trichotomous: treatment, control, or neither.



**Figure 3. Treatment classification comparison, Afrobarometer household survey locations**



Notes: The HP19 network maps are drawn in black, and national borders in grey. Hollow circles show Afrobarometer enumeration areas where the HP19 and new treatment classifications agree on treatment assignment under the HP19 classification (within  $0.005^\circ$  of a fiber line for treatment). Solid circles show disagreements. Circles are plotted at their AidData coordinates (BenYishay et al. 2017). Scales are chosen separately for each axis of each plot.

### 3.2 Replication and revision: Internet speed and usage

The first impact estimates in HP19 are for Internet speed and frequency of Internet use. The speed data come from Akamai and the usage data from Afrobarometer. In the language of instrumental variables, these regressions can be viewed as the first stage in HP19's assessment of the impact of submarine cables. They assess the association between the variable that is the primary source of identifying variation and more proximate indicators of treatment.

Unfortunately, the data analyzed in this first stage is among the most coarsely geocoded in the study. As a result, it appears, the estimates are fragile. The Akamai Internet speed data are coded only by city (or province or state in some cases). And for this first-stage analysis, HP19 replaces Afrobarometer

survey locations with the nearest Akamai cities. Notably, precisely because in the separate WBES firm survey locations are also only coded by city, HP19 deemphasizes results from that data set, reporting them in an appendix.<sup>18</sup>

Here, because the details of the mapping between Afrobarometer and Akamai locations are absent from the HP19 code archive, I reimplement the Internet usage regressions by copying HP19's Afrobarometer employment specifications. This change has the additional benefits of constituting a minimally arbitrary robustness test and more directly examining whether, within the analytical frame for estimating reduced-form employment impacts, first-stage effects can be detected.

The first three columns of the first row of Table 2 replicate HP19's estimates of the impact of submarine cable connection on Akamai-measured bandwidth. The result in the upper left corner repeats the original finding that undersea connection increased bandwidth by 0.35 arc-hyperbolic-sine ( $\text{asinh}$ ) points. That specification controls for country-year and location fixed effects and excludes location-quarters with 10 or fewer speed measurements. Following HP19, column 2 introduces an exclusion for large cities, since they ambiguously encompass neighborhoods within and beyond the reach of broadband. Also following HP19, column 3 adds treatment group-year dummies.

The final columns of the first row present the impact estimates for Afrobarometer-reported Internet usage by individuals—in particular, for whether people log on at least daily or at least weekly. Like the employment regressions reported in columns 5 and 7 of HP19 Table 2, the Internet usage regressions here control for country-year and treatment-year fixed effects, and do not exclude the

<sup>18</sup> A more precise reading of HP19 may be that it deemphasizes the WBES results for two reasons: “In light of the WBES firms being clustered in space *and* the approximate connectivity classification required due to the relatively aggregate geographical level at which their location is reported... [emphasis added].” However, all the survey data in HP19 comes from complex surveys that cluster the sampling, so the first reason does not seem like an operative distinction.

biggest cities. The point estimates are negative.

Much as in the previous table, the remaining rows of Table 2 accumulate specification changes. In the second row, HP19's definitions of the treatment and control zones are retained. But when working with the Akamai data whose cadence is quarterly, quarterly dummies replace yearly ones, in conformity with the TWFE design and the precedent set by HP19's regressions on South Africa QLFS data. The list of large cities is revised, as described in footnote 7. And for both Akamai and Afrobarometer, the revised geocodings are introduced. All full-sample impacts are now difficult to distinguish from zero; but the Akamai regressions excluding the largest cities are less affected. The third and fourth rows switch to the preferred treatment definitions, using radii of  $0.1^\circ$  and  $0.2^\circ$  from fiber lines or nodes. The estimates are usually negative, albeit without much statistical significance.

These null, even negative, results stand in tension with the clear treatment-control contrasts drawn in Table 1 when using  $0.1^\circ/0.2^\circ$  bounds. How are they squared? The earlier results are (post-treatment) differences, whereas those in Table 2 are differences in differences. The Table 1 results are for broadband proxies with a resolution of about 600 meters (for Ookla speeds) or 260 meters (for 3G coverage) whereas the Table 2 results reference Akamai download data points geolocated only to the city. As a result, while there is little doubt that proximity to fiber is associated with broadband access, and little doubt that undersea cable connection circa 2010 increased broadband availability in Africa, it is less clear that it did so *differentially* with respect to distance from fiber, at least enough so to generate a detectable signal in the specifications and data brought to bear here.

**Table 2. Impact of submarine cable arrival on Internet speed and usage**

Specification	Akamai			Afrobarometer	
	Log bandwidth			Use Internet daily	Use Internet weekly
Original, 0.005°/0.1° from lines	0.354 (0.134)	0.362 (0.171)	0.380 (0.185)	−0.016 (0.018)	−0.005 (0.012)
Observations	2,533	1,670	1,670	10,588	10,588
+ Re-geocode, recalculate distance	0.027 (0.148)	0.253 (0.166)	0.184 (0.233)	−0.019 (0.021)	0.029 (0.021)
Observations	2,499	1,621	1,621	10,211	10,211
+ Change to 0.1°/0.2° from lines	−0.216 (0.330)	−0.257 (0.466)	0.453 (0.251)	−0.044 (0.038)	−0.039 (0.025)
Observations	2,704	1,694	1,686	12,103	12,103
+ Change to 0.1°/0.2° from nodes	−0.156 (0.286)	−0.300 (0.321)	−0.094 (0.364)	−0.023 (0.025)	−0.009 (0.020)
Observations	2,837	1,837	1,837	10,532	10,532
Treatment-time FE	No	No	Yes	Yes	Yes
Include biggest cities	Yes	No	No	Yes	Yes

Note: Each cell holds an estimate of a coefficient on the treatment term, a dummy for being close to national fiber network after it has been connected to an undersea cable circa 2010. Bandwidth results in first row match corresponding values in HP19 Table 2. Internet usage results in rest of first row do not; their specifications match the HP19 employment regressions except in the dependent variables. For all first-row estimates, the treatment group is within 0.005° of a fiber line and the comparators within 0.005–0.1°. Subsequent rows accumulate changes toward the preferred specifications at bottom. In the second row, observations are re-geocoded, the list of four biggest cities in each country is revised, quarterly dummies replace yearly ones, and fixed-effect singletons are dropped. For the Internet usage results, a seemingly erroneous HP19 sample restriction is also dropped; see notes to Table 4. In the third row, bounds for the treatment and non-treatment groups are extended to 0.1° and 0.2°. In the last, fiber optic nodes become the referent for the geographic dimension of the treatment definition. All regressions include locality dummies. Standard errors clustered by locality are in parenthesis.

### 3.3 Replication and revision: Employment

HP19's primary outcome is employment, at the level of the individual or the firm. Impacts are estimated for overall employment and, as data allow, by the skill level of the work and the education level of the worker. Table 3, Table 4, and Table 5 revisit the individual-level employment results from the DHS, Afrobarometer, and South Africa QLFS data sets. Table 6 does the same for firm-level employment in the Ethiopia LMMIS. All follow Table 2 in format.

## DHS

The first major row of Table 3 recapitulates the HP19 employment results from DHS data: submarine connection is followed by a 4.6 percentage point increase in employment, with the effect concentrated in skilled occupations but spread rather evenly across education groups. All the underlying regressions include controls for country-time effects and grid cell–treatment status effects, where grid cells are  $0.1^\circ \times 0.1^\circ$ . Standard errors are clustered by grid cell. The samples cover eight countries, with one pre- and one post-treatment survey from each.

The second major row sticks to the HP19  $0.005^\circ/0.1^\circ$ -from-lines treatment classification but introduces a correction. An erroneous exclusion in the HP19 code of observations whose interview year differs from the nominal survey year is reversed; this matters for DHS surveys that interviewed people across two calendar years. The sample expansion lowers the overall employment impact by a third and removes the impression that the effect is biased toward skilled work.

Drawing treatment and control boundaries at  $0.1^\circ$  and  $0.2^\circ$ , whether from lines or nodes, further reduces the apparent impact on overall employment (last two major rows of Table 3). To complement the geographic coarsening of the treatment and comparator definitions, grid cells are now enlarged to  $0.2^\circ \times 0.2^\circ$ , as in Hjort and Poulsen (2016). The overall employment estimate (first column) is 1.5 points when taking distance from lines (standard error 1.4) and 2.3 points (1.4), when taking distance from nodes. To this extent, the HP19 finding of an employment impact in the DHS data is preserved. The finding of skill bias looks less robust.

**Table 3. Impact of submarine cable arrival on individual employment (DHS)**

Specification	Any employment	Skill level of job		Education level			
		Skilled	Unskilled	Primary not completed	Primary	Secondary	Higher
Original, 0.005°/0.1° from lines	0.046 (0.014)	0.038 (0.018)	0.004 (0.014)	-0.013 (0.032)	0.061 (0.022)	0.047 (0.015)	0.067 (0.027)
Observations	59,914	59,914	59,914	59,912	59,912	59,912	59,912
+ Revise sample	0.030 (0.013)	0.013 (0.015)	0.018 (0.011)	-0.016 (0.021)	0.037 (0.018)	0.030 (0.014)	0.043 (0.022)
Observations	123,114	123,114	123,114	123,102	123,102	123,102	123,102
+ Change to 0.1°/0.2° from lines	0.015 (0.014)	0.006 (0.014)	0.008 (0.016)	-0.002 (0.018)	0.031 (0.015)	0.013 (0.015)	-0.014 (0.018)
Observations	159,408	159,408	159,408	159,394	159,394	159,394	159,394
+ Change to 0.1°/0.2° from nodes	0.023 (0.014)	-0.009 (0.012)	0.032 (0.015)	0.005 (0.020)	0.038 (0.016)	0.025 (0.014)	-0.011 (0.018)
Observations	131,387	131,387	131,387	131,376	131,376	131,376	131,376

Note: Results in first row exactly or closely match those in HP19 Table 3, panel A, column 1; Table 5, panel A, columns 1 and 2; and Table 6, first panel, column 1. For them, treated units lie within 0.005° of a fiber line and comparators within 0.005–0.1°. Subsequent rows accumulate changes. In the second, an erroneous exclusion of observations with interview year different from nominal survey year is reversed. In the third, treatment and comparator radii are extended to 0.1° and 0.2°. In the last, distance is taken from the nearest network node rather than network line, using 2010 node locations from Africa Bandwidth Maps. In each row, results in the last four columns are coefficients on interaction terms within the same regression. All regressions include dummies for country-year and grid cell-treatment group combinations. Grid cells are 0.1°×0.1° in the first two rows, as in HP19, and 0.2°×0.2° in the last two, as in Hjort and Poulsen (2016). Standard errors, clustered by 0.1°×0.1° grid cell in the first two rows and 0.2°×0.2° in the last two, are in parenthesis.

### Afrobarometer

The first row of the Afrobarometer impacts table (Table 4) replicates the HP19 estimates for total employment and employment by education level. In the second row, a handful of observations with interview year differing from nominal survey year are restored. Also reversed is a sample exclusion unique within HP19, of people who are neither working nor looking for work; they are now counted as not employed.<sup>19</sup> And the new geocoding is introduced. These

<sup>19</sup> The HP19 exclusion would also be feasible, but is not applied, for the South Africa QLFS data since that survey also asks whether people were recently looking for work. Retaining the HP19 exclusion for the Afrobarometer data changes the

revisions erase the apparent impact on employment in this data set. Widening the treatment and control zones, as reported in the third and fourth major rows, also leads to point estimates hard to distinguish from zero.

**Table 4. Impact of submarine cable arrival on individual employment (Afrobarometer)**

Specification	Any employment	Education level			
		Primary not completed	Primary	Secondary	Higher
Original, 0.005°/0.1° from lines	0.077 (0.036)	0.113 (0.053)	0.081 (0.040)	0.066 (0.048)	0.105 (0.051)
Observations	7,918	7,862	7,862	7,862	7,862
+ Revise sample, recalculate distance	0.006 (0.036)	−0.021 (0.050)	−0.030 (0.040)	0.010 (0.044)	0.013 (0.051)
Observations	10,179	10,162	10,162	10,162	10,162
+ Change to 0.1°/0.2° from lines	0.038 (0.046)	0.071 (0.048)	0.030 (0.046)	0.060 (0.050)	0.035 (0.054)
Observations	12,063	12,045	12,045	12,045	12,045
+ Change to 0.1°/0.2° from nodes	−0.026 (0.034)	0.019 (0.039)	−0.032 (0.036)	−0.019 (0.037)	−0.047 (0.043)
Observations	10,494	10,477	10,477	10,477	10,477

Note: This table has the same format as Table 3 except for lacking columns for skilled and unskilled employment. See notes to that table. Results in first row exactly or closely match those in HP19 Table 3, panel A, column 2, and Table 6, second panel, column 1. Starting in the second row, regressions drop an exclusion of unemployed working-age adults who are not looking for work. Standard errors, clustered by 0.1°×0.1° grid cell in the first two rows and 0.2°×0.2° in the last two, are in parenthesis.

### South Africa QLFS

Results for the South Africa labor force survey appear in Table 5. Although the data are technically re-geocoded starting in the second row—the enumeration area centroids are recomputed—the treatment-control classification hardly changes (see Figure A–5). Following HP19, enumeration areas replace grid cells as the basis for spatial fixed effects and errors clustering. Where HP19 finds an employment boost of 2.2 points in South Africa, the wider-zone specifications, in the bottom half of Table 5, find none. And, as in the new DHS results, there is no clear bias for or against skilled employment.

estimated employment impact of 0.006 (standard error 0.036) in row 2 of Table 4 to −0.008 (0.038).

**Table 5. Impact of submarine cable arrival on individual employment (South Africa QLFS)**

Specification	Any em- ployment	Skilled	Unskilled	Primary not completed	Primary	Secondary	Higher
Original, 0.005°/ 0.1° from lines	0.022 (0.008)	0.012 (0.009)	0.010 (0.007)	0.012 (0.017)	0.028 (0.011)	0.022 (0.012)	0.019 (0.011)
Observations	280,641	280,641	280,641	277,737	277,737	277,737	277,737
+ Recalculate distance	0.025 (0.009)	0.018 (0.010)	0.007 (0.008)	0.019 (0.020)	0.034 (0.013)	0.023 (0.013)	0.017 (0.013)
Observations	270,422	270,422	270,422	267,600	267,600	267,600	267,600
+ Change to 0.1°/0.2° from lines	-0.004 (0.010)	0.004 (0.010)	-0.008 (0.008)	-0.011 (0.011)	-0.011 (0.010)	-0.003 (0.010)	0.003 (0.010)
Observations	307,187	307,187	307,187	304,037	304,037	304,037	304,037
+ Change to 0.1°/0.2° from nodes	0.000 (0.006)	0.002 (0.006)	-0.002 (0.006)	-0.002 (0.010)	-0.006 (0.007)	0.005 (0.008)	0.006 (0.008)
Observations	280,609	280,609	280,609	277,686	277,686	277,686	277,686

Note: This table has the same format as Table 3, which see. Results in first row exactly or closely match those in HP19 Table 3, panel A, column 3; Table 5, panel B, columns 1 and 2; and Table 6, third panel, column 1. Unlike in the previous tables, grid cell fixed effects are not controlled for, and all standard errors are clustered by enumeration area.

### Ethiopia LMMIS

HP19 finds a large effect on firm-level employment in Ethiopia, at 0.156 or 0.224 log points depending on the choice of fixed effects. Here too, skilled employment appears to increase more. Those results appear in the top row of Table 6. The revision to the geocoding of Ethiopia LMMIS observations documented in Figure A–6 hardly does not appreciably change these estimates (second row). However, as in South Africa, revising the treatment and control zones greatly weakens the results. Then, almost none are easily distinguishable from zero.



**Table 6. Impact of submarine cable arrival on firm employment (Ethiopia LMMIS)**

Specification	Employees		Skilled employees		Unskilled employees	
Original, 0.005°/ 0.1° from lines	0.156 (0.091)	0.224 (0.081)	0.034 (0.080)	0.201 (0.106)	0.115 (0.078)	0.124 (0.095)
Observations	5,360	5,360	5,360	5,360	5,360	5,360
+ Re-geocode, recalculate distance	0.151 (0.083)	0.197 (0.067)	0.044 (0.105)	0.222 (0.095)	0.259 (0.063)	0.172 (0.087)
Observations	5,737	5,696	5,737	5,696	5,737	5,696
+ Change to 0.1°/0.2° from lines	-0.092 (0.123)	0.018 (0.134)	-0.055 (0.267)	-0.334 (0.378)	-0.309 (0.106)	-0.065 (0.099)
Observations	5,748	5,707	5,748	5,707	5,748	5,707
+ Change to 0.1°/0.2° from nodes	-0.042 (0.105)	0.040 (0.148)	-0.082 (0.144)	0.017 (0.179)	-0.191 (0.108)	-0.058 (0.166)
Observations	5,269	5,228	5,269	5,228	5,269	5,228
Year FE	Yes	No	Yes	No	Yes	No
Grid cell-connected FE	No	Yes	No	Yes	No	Yes
Industry-year FE	No	Yes	No	Yes	No	Yes
Firm FE	Yes	No	Yes	No	Yes	No

Note: This table has the same format as Table 3, which see. Results in first row exactly or closely match those in HP19 Table 8, first row. Dependent variables are subject to the arc-hyperbolic sine transform. Standard errors, clustered by 0.1°×0.1° grid cell in the first two rows and 0.2°×0.2° in the last two, are in parenthesis.

### World Bank Enterprise Surveys

HP19 runs similar specifications on the firm-level data from the multi-country WBES in order to study impacts on firm-level exports, on-the-job training, and corporate use of the Internet. The regressions indicate a 14–17 log point increase in firm employment after undersea cable connection. But, as noted above, HP19 places less confidence in the results because survey locations are only coded by city.

HP19's deemphasis of the WBES results is even more warranted than it appears at first. All of the identification flows from Nigeria data, where the geographic coding is especially coarse. Concretely, if Nigeria is dropped from the WBES regressions, the treatment indicator becomes collinear with the fixed-effect dummies. Why? In the HP19 subset of the WBES data, Nigeria is the only country to have more than four locales visited in any one round. In each of the

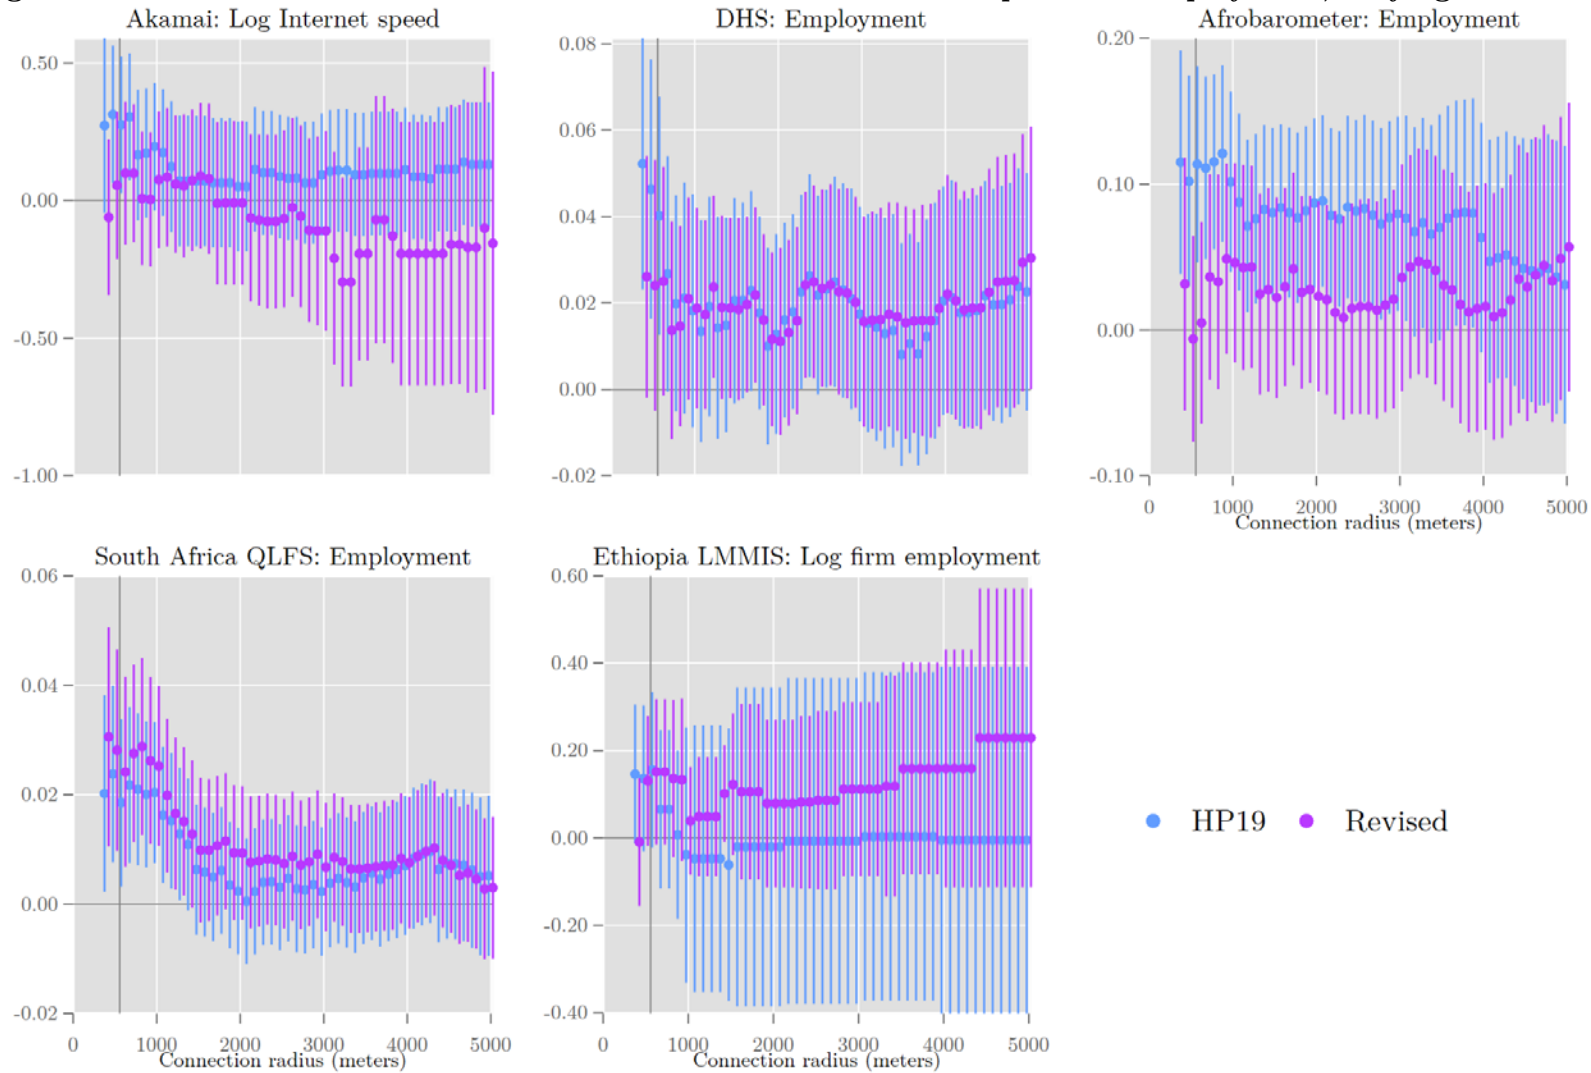
other countries, those cities that received surveyors both before and after submarine cable connection were either all deemed close to fiber lines in the HP19 coding, or all far. Only the Nigeria data can fill all the boxes in the classical  $2 \times 2$  difference-in-differences grid, and therefore contribute identifying information under HP19's pooled, multi-country TWFE design. And for Nigeria, the "firm establishment location" field turns out to contain names of states, not cities. Since firms cannot be geolocated with much precision when only their home states are known, I fully set aside the WBES analysis.

### 3.4 Varying the treatment radius more systematically

While HP19 prefers a treatment radius of  $0.005^\circ$  (~560 meters), the study performs a robustness check on that choice. In the preferred DHS, Afrobarometer, and South Africa QLFS specifications, Figure 4 of HP19 varies the radius between 400 and 2,000 meters for DHS and South Africa QLFS and between 400 and 5,000 meters for Afrobarometer, and displays the evolution of the impact estimates. To address the tension between the arguable robustness in that figure and the greater fragility reported here, I replicate and extend the HP19 figure. I add the Akamai and Ethiopia regressions to the testing, and bring Afrobarometer's radius range, 400–5,000 meters, to all datasets. I also incorporate the revisions embodied in the second row of the above tables, for example, including observations whose interview year differs from the nominal survey year. And for the multi-country specifications, I correct one error in the HP19 code: while the tabulated HP19 regression results incorporate country–time fixed effects, the Figure 4 regressions incorporate only time effects. The new results are plotted here in Figure 4.

The dominant theme in the results is that in the data sets where geocoding is revisited because the primary data contains only place names—Akamai, Afrobarometer, and Ethiopia LMMIS, the results change substantially. The Akamai coefficient is hard to distinguish from zero in both the original and revised specifications. The Afrobarometer coefficient is much lower in the revised

specification, while the Ethiopia coefficient is higher. Results are more stable in the DHS and South Africa QLFS data sets. In the DHS, the employment impact hovers around 2 percentage points, which coheres with the preferred DHS estimates in the lower half of Table 3 above. In the South Africa data, the estimate moves closer to zero as the treatment radius extends to 2,000 meters, stabilizes at roughly 1 point till about 4,000 meters, and then declines further. Since HP19 does not provide a specific technological basis for the  $0.005^\circ$  radius, and since the evidence in Table 1 favors longer radii, it is reasonable to put less interpretive weight on the results at the left edges of these plots, though they correspond to HP19's preferred estimates.

**Figure 4. Estimated effect of submarine cable connection on internet speed and employment, varying the treatment radius**

Notes: Following HP19, Figure 4, each plot shows how the results in the first columns, first and second rows, of Table 2–Table 6 vary as the outer bound for the treatment zone is varied between 400 and 5,000 meters.

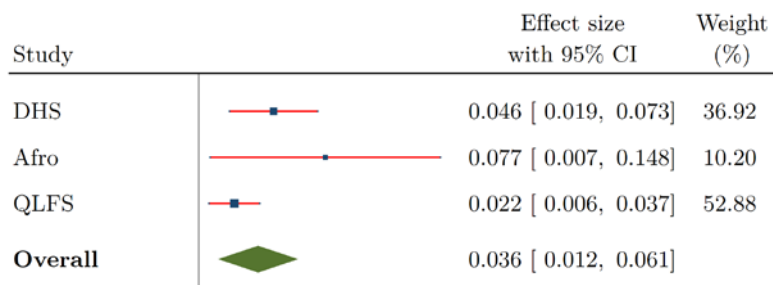
### 3.5 Meta-analysis

To distill the results on individual-level employment from the DHS, Afrobarometer, and South Africa QLFS data sets, I perform meta-analysis. The spirit of the exercise is to average the estimates in a way that factors in their relative precision and minimizes researcher discretion.<sup>20</sup> The default method in Stata’s “meta forestplot” command, DerSimonian-Laird random-effects meta-analysis, is applied. This method does not assume that the average treatment effect is homogeneous across settings, and here must estimate the heterogeneity with a sample of only three “studies.” See Figure 5. The meta-analytically averaged estimate of the impact of broadband on employment is 3.6 percentage points for the HP19 data and specifications, with a standard error of 1.2 points (first pane). The estimate drops to 2.6 points (standard error 0.7) with data, specification, and geocoding revisions (second pane), and to 0.4 or 0.5 points (standard error 0.9 for both) when classifying with 0.1°/0.2°-from-lines or 0.1°/0.2°-from-nodes (last two panes). The main source of positivity in the revised results is the DHS data, where the lag between pre- and post-treatment measurement is longest (9 years), and where, as a result, the exogenous shocks of submarine cable connection account for the smallest share of the temporal component of identifying variation.

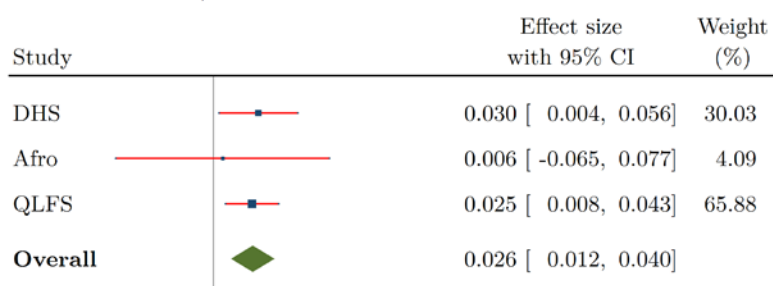
<sup>20</sup> Otis Reid suggested this step.

**Figure 5. Meta-analytical summary: fast Internet and employment**

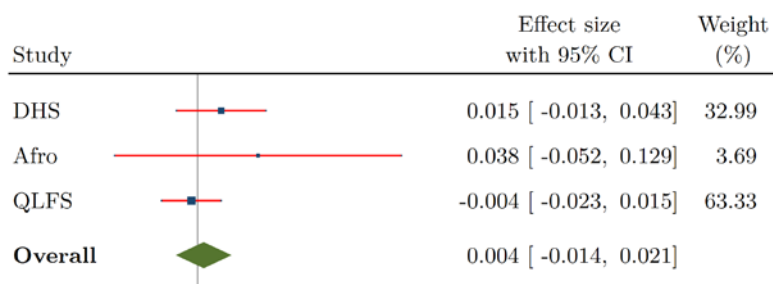
Original, 0.005°/0.1° from lines



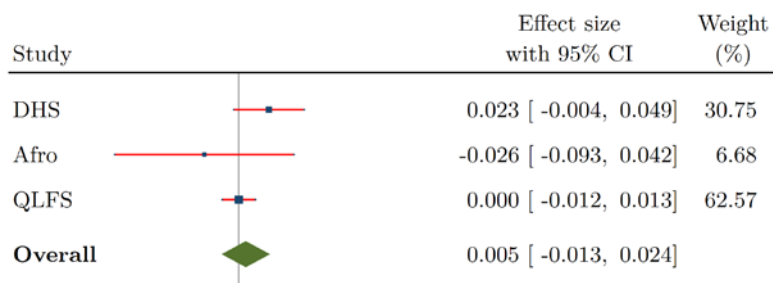
Revised, 0.005°/0.1° from lines



Revised, 0.1°/0.2° from lines



Revised, 0.1°/0.2° from nodes



Notes: These forest plots combine employment impact estimates from the first columns of Table 3, Table 4, and Table 5 using DerSimonian and Laird random-effects meta-analysis.

## 4 Results for nighttime light

Analysis of satellite data is a promising way to quantify economic development (Elvidge et al. 1997; Chen and Nordhaus 2011; Henderson, Storeygard, and Weil 2012). But there are pitfalls, especially when using data from the early-generation DMSP-OLS satellite system. That system was designed not for social science but for measuring cloud cover to help the U.S. Air Force forecast the weather (Gibson, Olivia, and Boe-Gibson 2020).

One major complication in the “stable lights” DMSP data sets used in economics is that the values are not strictly comparable over time. The Air Force regularly adjusted the satellite sensors, notably to compensate for the waxing and waning of the moon. The specifics of this variation are not documented. Moreover, the stable light values are annual averages at each location over cloud-free days. Since cloud cover variability could intersect with the lunar cycle in complex ways, the amplification baked into each pixel varies over space and time. A final complication is that a major disruption occurred within the HP19 timeframe: the DMSP switched satellites between 2009 and 2010.<sup>21</sup> At that juncture, the global sum of the light data jumped by more than half—by far the largest year-to-year change (Li and Zhou 2017, Figure 1).

These problems are well recognized, at least outside of economics (Gibson, Olivia, and Boe-Gibson 2020). And perhaps they are why, within economics, researchers tend to use the DMSP data in ways that minimize the influence of short-term changes, working with pure cross-sections (Bleakley and Lin 2012, Michalopoulos and Papaioannou 2013) or long differences.<sup>22</sup>

A common response to the incomparability over time is to “intercalibrate” different years’ data using the observations of *Sicily* (Elvidge et al. 2009). Since

<sup>21</sup> Earth Observation Group, [eogdata.mines.edu/products/dmsp/#download](http://eogdata.mines.edu/products/dmsp/#download), accessed April 8, 2023.

<sup>22</sup> “Most empirical studies using light data have focused on cross-section variation.” Lowe (2014, p.10).

Sicily is part of a mature economy, its light emissions might not change much year-to-year. And its pixel-level readings are relatively evenly distributed across the measured power range. Typically, to calibrate against Sicily, the coefficients in a quadratic transformation are tuned to minimize the least-squares discrepancy between the benchmark year's readings for the territory and a given year's intercalibrated values. The quadratic transformation is then applied globally.

HP19 takes an alternative approach: forgoing calibration and relying on year dummies to expunge systematic shifts over time, as in Chen and Nordhaus (2011) and Henderson, Storeygard, and Weil (2012).

Neither approach works perfectly. As just explained, whatever causes the gain shifts does not affect all pixels equally. What's best for Sicily is not best for Africa. And unremoved artifacts that are second-order when taking first differences, such as when comparing regions, can become first-order in second differences. For its part, the fixed-effect approach only works to the extent that the gain shifts respect the transformations that the researcher applies to the light data. In the best case, which is unrealistic, researchers take the light values in logarithms, and calibrating each year's data only requires multiplication by some global constant. Then, the gain will be perfectly absorbed by year fixed effects. But if that were true, intercalibration wouldn't call for a quadratic transform. And in HP19, the light values enter regressions after passing through the arc-hyperbolic sine ( $\text{asinh}$ ) function. It would be surprising if, after applying this transform, the time effects are constant across the power spectrum, thus fully expungeable by fixed-effect controls.

I replicate and revise the HP19 nighttime light estimates following the template of the previous section. In addition, I incorporate three distinctive changes: fixing a coding error that causes some countries' post-treatment dummies to exceed 1; working with the data at their original resolution of  $1/120^\circ$  instead of  $1/10^\circ$ , in order to assign treatment status more precisely; and switching to the intercalibrated version of the DMSP data from the Earth Observation Group,



which is also the primary source of the uncalibrated data sets.<sup>23,24</sup> As with the non-satellite estimates, HP19 controls for country-year and grid cell-treatment zone fixed effects in both its light regression. In the second regression, HP19 also adds treatment zone-year effects. These additional, time-varying controls look especially valuable in view of the concern just voiced about temporal artifacts in the data.

The HP19 result without those extra fixed effects proves robust, as shown in the first and third columns of Table 7. The specification produces large estimates even when using the higher-resolution, intercalibrated data and the node-based treatment classification. Indeed, the point estimate of the impact on  $\text{asinh}(\text{light})$  more than doubles, from 0.024 in the original to 0.053 (fourth row, third column). A major reason appears to be the correction of the coding error: shrinking the range of the treatment indicator to a strict  $[0,1]$  naturally increases its coefficient.

However, copying HP19 in adding treatment zone-year fixed effects greatly reduces point estimates (columns 2 and 4 of Table 7). To gain insight into why, I plot the estimates of these added fixed effects, taking 2013 as the reference year. I do so in regressions corresponding to the bottom half of Table 7, with both raw and intercalibrated data at original resolution. I estimate in the full

<sup>23</sup> Many authors have developed intercalibrated versions of the DMSP-OLS data (Li et al. 2013; Wu et al. 2013; Hsu et al. 2015; Stathakis 2006; Li and Zhou 2017; Zhao et al. 2022). The minimally arbitrary choice is to stick to the Earth Observation Group (Hsu et al. 2015) since it is the sole source of the non-intercalibrated data.

<sup>24</sup> Initially, I also fit a Poisson model to the untransformed light values (Correia, Guimarães, and Zylkin 2020). However, these results are sensitive to the coding of the darkest pixels. While the documentation for the light data states that values run between 1 and 63, the darkest pixels are coded with 0, and 1 never occurs. The Poisson results are much more sensitive than the linear results to whether dark pixels are coded with 0 or 1. Were it computationally feasible, a fixed-effect, bi-censored tobit model might be more appropriate for untransformed light values, since the sensors fail to detect signals below some threshold and saturate at the high end.

sample, and in each country sample separately. See Figure 6. In the raw light data, even though different countries received undersea cable connection in different years, nearly all countries' treatment zone effects jump with statistical significance between 2009 and 2010, when the DMSP data switches from satellite F16 to F18. Instead estimating on the intercalibrated light data (right half of the figure) reduces but does not eliminate the shared jump. In the fiber line-based classification (upper right of the figure) all countries' series still rise at the satellite switch, most with great statistical significance. The same appears to hold under the node-based classification (lower-right), except that the estimates are noisier, making the rises less statistically significant.

Since the treatment zone-year interactions are probably picking up the satellite switch, it makes sense to rely on results from the specification that control for those dummies. These are smaller and appear in the even columns of Table 7. At the same time, the right half of Figure 6 suggests that the bias from the satellite switch lingers even after including these fixed effects.

**Table 7. Impact of submarine cable arrival on nighttime light**

	Raw light values		Intercalibrated light values	
Original, 0.005°/0.1°	0.024	0.033		
from lines	(0.009)	(0.017)		
Observations	80,360			
+ Recalculated distance,	0.046	0.018	0.037	0.016
higher resolution	(0.003)	(0.004)	(0.003)	(0.004)
Observations	11,530,295		11,519,153	
+ Change to 0.1°/0.2°	0.038	0.011	0.030	0.012
from lines	(0.002)	(0.003)	(0.002)	(0.003)
Observations	20,093,941		20,076,115	
+ Change to 0.1°/0.2°	0.069	0.002	0.050	0.006
from nodes	(0.004)	(0.006)	(0.004)	(0.006)
Observations	9,939,804		9,930,294	
Country-year FE	Yes	Yes	Yes	Yes
Grid cell-connected FE	Yes	Yes	Yes	Yes
Connected-year FE	No	Yes	No	Yes

Note: This table has the same format as Table 3. See notes to that table. Light values, ranging between 0 and 63, are linearly modeled after taking their arc-hyperbolic sines. Results in first row nearly match those in HP19 Table 9. Starting in the second row, the post-treatment indicator does not exceed 1 in countries that received cable multiple connections, correcting an error in the original. Raw and intertemporally calibrated values are from the Earth Observation Group. Standard errors, clustered by 0.1°×0.1° grid cell in the first two rows and 0.2°×0.2° in the last two, are in parenthesis.

**Figure 6. Evolution of association between stable nighttime light and proximity to fiber network, by country**

Notes: Each plot shows the estimated coefficients on products of the connected dummy and year dummy, along with 95% confidence intervals, in regressions whose specifications match those in the bottom two rows, second and fourth columns, of Table 7. The regressions are run separately for each country in the HP19 light sample, and for all at once. The Defense Meteorological Satellite Program-Operational Linescan System stable light data set switches from satellite F16 to F18 between 2009 and 2010.

## 5 Conclusion

Reanalysis of HP19 weakens the evidence that undersea cable connections circa 2010 stimulated economic activity in Africa, in particular, increasing employment with a positive skill bias.

The largest reason for this revised view is the lack of robustness to changing the treatment and comparator zone radii from  $0.005^\circ/0.1^\circ$  to  $0.1^\circ/0.2^\circ$ . No definition of treatment and control zones can be conclusively preferred. But the new choices have several virtues. They are minimally arbitrary, as round numbers that define zones of equal width ( $0.1^\circ$ ). They have demonstrably higher fidelity to high-resolution indicators of broadband availability, and so should generate more statistical power. And they are more compatible with the resolution of the various geographic measurements used here. In other words, they step away from the pairing in HP19 of stronger results with weaker identification. Under the new radii, the meta-analytically averaged estimates of individual-level employment impacts are hard to distinguish from zero, at 0.4 or 0.5 points (standard error 0.9). The slight positivity in this bottom line derives from regressions on the samples with the longest lags between pre- and post-treatment measurement—9 years, in the DHS data—thus the weakest claim to causal identification. Evidence of positive impact on firm-level employment disappears in the Ethiopia LMMIS surveys, as well as the WBES ones. Signs of impact on light emissions, as a proxy for economic activity, look mostly like artifacts of a satellite changeover. Throughout, the signs of skill bias also disappear.

Perhaps these results should not surprise. Economists have long rebutted the Luddite fallacy that technological advance destroys jobs. Generalizing, a reasonable prior is that any given technology's effects on aggregate employment are small and hard to detect.

## References

- Bahia, Calvin, Pau Castells, Carlos Rodríguez-Castelán, Takaaki Masaki, Genaro Cruz, Tobias Pfutze, Hernan Winkler, and Xavier Pedrós. 2020. "The Welfare Effects of Mobile Broadband Internet: Evidence from Nigeria." <https://doi.org/10.2139/ssrn.3757666>.
- BenYishay, Ariel, Renee Rotberg, Jessica Wells, Zhonghui Lv, Seth Goodman, Lidia Kovacevic, and Dan Runfolo. 2017. "Geocoding Afrobarometer Rounds 1–6: Methodology & Data Quality." AidData. <https://docs.aiddata.org/ad4/pdfs/geocodingafrobarometer.pdf>.
- Bischof, Daniel. 2017. "New Graphic Schemes for Stata: plotplain and plottig." *Stata Journal* 17 (3): 748–59. [doi.org/10.1177/1536867X1701700313](https://doi.org/10.1177/1536867X1701700313).
- Chen, Xi, and William D. Nordhaus. 2011. "Using Luminosity Data as a Proxy for Economic Statistics." *Proceedings of the National Academy of Sciences of the United States of America* 108 (21): 8589–94. <https://doi.org/10.1073/pnas.1017031108>.
- Chiplunkar, Gaurav, and Pinelopi Koujianou Goldberg. 2022. "The Employment Effects of Mobile Internet in Developing Countries." Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w30741>.
- Correia, Sergio. 2014. "REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects." Revised 18 Nov 2019. <https://ideas.repec.org/c/boc/bo-code/s457874.html>.
- Correia, Sergio, Paulo Guimarães, and Tom Zylkin. 2020. "Fast Poisson Estimation with High-Dimensional Fixed Effects." *Stata Journal* 20 (1): 95–115. <https://doi.org/10.1177/1536867X20909691>.
- Elvidge, Christopher D., Kimberly E. Baugh, Eric A. Kihn, Herbert W. Kroehl, and Ethan R. Davis. 1997. "Mapping City Lights with Nighttime Data from the DMSP Operational Linescan System." *Photogrammetric Engineering and Remote Sensing* 63 (6): 727–34. <https://asprs.org/wp->

- [content/uploads/pers/97journal/june/1997\\_jun\\_727-734.pdf](content/uploads/pers/97journal/june/1997_jun_727-734.pdf).
- Elvidge, Christopher, Jeffrey Safran, Ingrid L. Nelson, Benjamin T. Tuttle, Vinita Ruth Hobson, Kimberly E. Baugh, John B. Dietz, and Edward H. Erwin. 2004. "Area and Positional Accuracy of DMSP Nighttime Lights Data." Ross S. Lunetta and John G. Lyon, eds. *Remote Sensing and GIS Accuracy Assessment*. CRC Press.
- Elvidge, Christopher D., Daniel Ziskin, Kimberly E. Baugh, Benjamin T. Tuttle, Tilottama Ghosh, Dee W. Pack, Edward H. Erwin, and Mikhail Zhizhin. 2009. "A Fifteen Year Record of Global Natural Gas Flaring Derived from Satellite Data." *Energies* 2 (3): 595–622.  
<https://doi.org/10.3390/en20300595>.
- Gibson, John, Susan Olivia, and Geua Boe-Gibson. 2020. "Night Lights in Economics: Sources and Uses." *Journal of Economic Surveys* 34 (5): 955–80.  
<https://doi.org/10.1111/joes.12387>.
- Goldbeck, Moritz, and Valentin Lindlacher. 2022. "Digital Infrastructure and Local Economic Growth: Early Internet in Sub-Saharan Africa."  
[https://moritz.goldbeck.net/documents/africa\\_broadband.pdf](https://moritz.goldbeck.net/documents/africa_broadband.pdf).
- Gonzalez, Robert M. 2021. "Cell Phone Access and Election Fraud: Evidence from a Spatial Regression Discontinuity Design in Afghanistan." *American Economic Journal. Applied Economics* 13 (2): 1–51.  
<https://doi.org/10.1257/app.20190443>.
- Guriev, Sergei, Nikita Melnikov, and Ekaterina Zhuravskaya. 2021. "3G Internet and Confidence in Government." *The Quarterly Journal of Economics* 136 (4): 2533–2613. <https://doi.org/10.1093/qje/qjaa040>.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2012. "Measuring Economic Growth from Outer Space." *American Economic Review* 102 (2): 994–1028. <https://doi.org/10.1257/aer.102.2.994>.
- Hjort, Jonas, and Jonas Poulsen. 2016. "The Arrival of Fast Internet and Skilled Job Creation in Africa."

- <https://www0.gsb.columbia.edu/mygsb/faculty/research/pub-files/19262/HjortPoulsen-AfricaInternet-091016.pdf>.
- Hjort, Jonas, and Jonas Poulsen. 2019. "The Arrival of Fast Internet and Employment in Africa." *American Economic Review* 109 (3): 1032–79. <https://doi.org/10.1257/aer.20161385>.
- Hsu, Feng-Chi, Kimberly E. Baugh, Tilottama Ghosh, Mikhail Zhizhin, and Christopher D. Elvidge. 2015. "DMSP-OLS Radiance Calibrated Nighttime Lights Time Series with Intercalibration." *Remote Sensing* 7 (2): 1855–76. <https://doi.org/10.3390/rs70201855>.
- Ioannidis, John P. A., T. D. Stanley, and Hristos Doucouliagos. 2016. "The Power of Bias in Economics Research." *Economic Journal* 127 (605): F236–F265. <https://doi.org/10.1111/ecoj.12461>.
- Jann, Ben. 2007. "Making Regression Tables Simplified." *Stata Journal* 7 (2): 227–44. <https://doi.org/10.1177/1536867X0700700207>.
- Klonner, Stefan, and Patrick J. Nolen. 2010. "Cell Phones and Rural Labor Markets: Evidence from South Africa." <https://www.econstor.eu/handle/10419/39968>.
- Li, Xi, Xiaoling Chen, Yousong Zhao, Jia Xu, Fengrui Chen, and Hui Li. 2013. "Automatic Intercalibration of Night-Time Light Imagery Using Robust Regression." *Remote Sensing Letters* 4 (1): 45–54. <https://doi.org/10.1080/2150704X.2012.687471>.
- Li, Xuecao, and Yuyu Zhou. 2017. "A Stepwise Calibration of Global DMSP/OLS Stable Nighttime Light Data (1992–2013)." *Remote Sensing* 9 (6): 637. <https://doi.org/10.3390/rs9060637>.
- Lowe, Matt. 2014. "Rail Revival in Africa? The Impact of Privatization." In *Joint RES-SPR Conference on Macroeconomic Challenges Facing Low-income Countries: New Perspectives*. <https://www.imf.org/external/np/seminars/eng/2014/lic/pdf/Lowe.pdf>.
- Manacorda, Marco, and Andrea Tesei. 2020. "Liberation Technology: Mobile

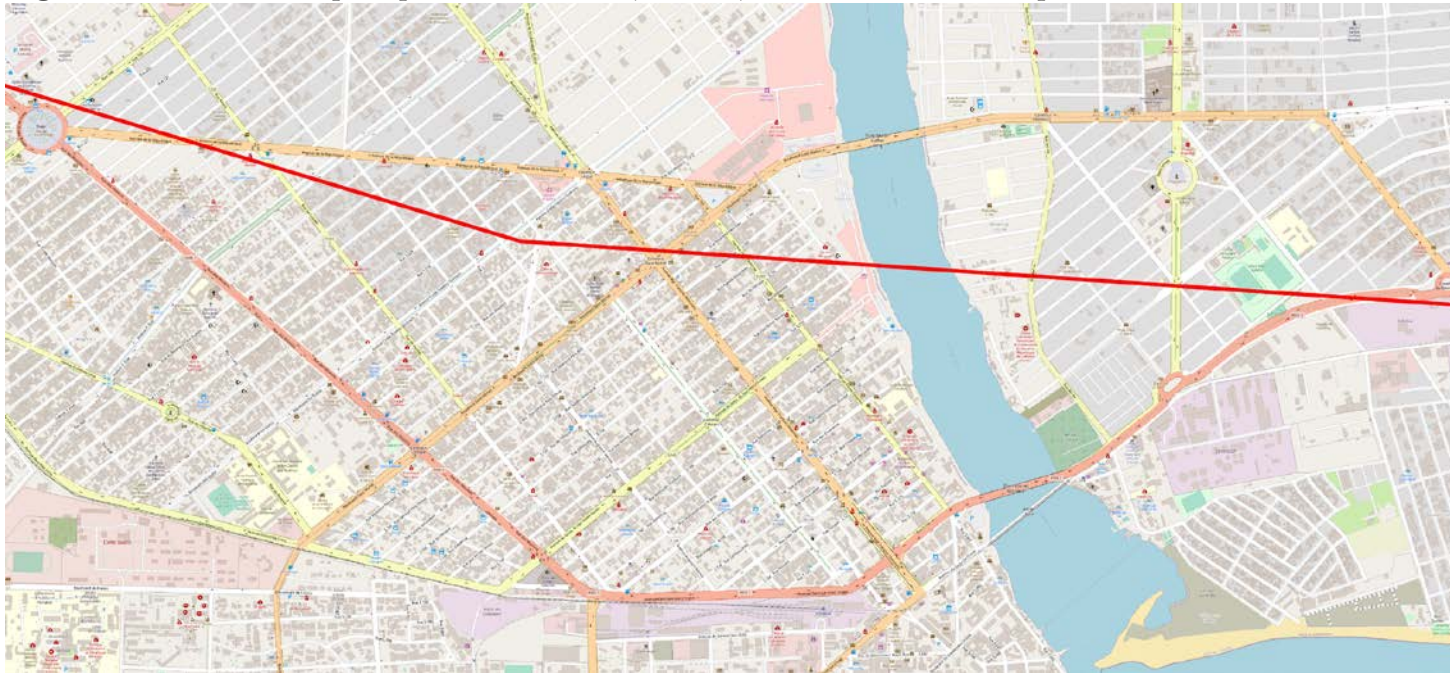


- Phones and Political Mobilization in Africa.” *Econometrica* 88 (2): 533–67.  
<https://doi.org/10.3982/ECTA14392>.
- Masaki, Takaaki, Rogelio Granguillhome Ochoa, and Carlos Rodríguez-Castelán.  
2020. “Broadband Internet and Household Welfare in Senegal.”  
<https://doi.org/10.1596/1813-9450-9386>.
- Mensah, Justice Tei, Kibrom Tafere, and Kibrom A. Abay. 2022. “Saving Lives  
through Technology: Mobile Phones and Infant Mortality.” Policy Research  
Working Papers. World Bank. <https://doi.org/10.1596/1813-9450-9978>.
- Stathakis, Demetris. 2016. “Intercalibration of DMSP/OLS by Parallel Regres-  
sions.” *IEEE Geoscience and Remote Sensing Letters* 13 (10): 1420–24.  
<https://doi.org/10.1109/LGRS.2016.2582784>.
- Over, Mead. 2022. “grc1leg2. Add a common legend to combined graphs.”  
<http://digital.cgdev.org/doc/stata/MO/Misc>.
- Perez-Heydrich, Joshua L. Warren, Clara R. Burget, and Michael E. Emch.  
2013. *Guidelines on the Use of DHS GPS Data*. [https://dhspro-  
gram.com/pubs/pdf/SAR8/SAR8.pdf](https://dhsprogram.com/pubs/pdf/SAR8/SAR8.pdf).
- Roodman, David. 2024. “Julia as a Universal Platform for Statistical Software  
Development.” arXiv [Econ.EM]. <http://arxiv.org/abs/2404.09309>.
- Wu, Jiansheng, Shengbin He, Jian Peng, Weifeng Li, and Xiaohong Zhong.  
2013. “Intercalibration of DMSP-OLS Night-Time Light Data by the Invari-  
ant Region Method.” *International Journal of Remote Sensing* 34 (20):  
7356–68. <https://doi.org/10.1080/01431161.2013.820365>.
- Zhao, Chenchen, Xin Cao, Xuehong Chen, and Xihong Cui. 2022. “A Consistent  
and Corrected Nighttime Light Dataset (CCNL 1992-2013) from DMSP-OLS  
Data.” *Scientific Data* 9 (1): 424. [https://doi.org/10.1038/s41597-022-01540-  
x](https://doi.org/10.1038/s41597-022-01540-x).

## Appendices

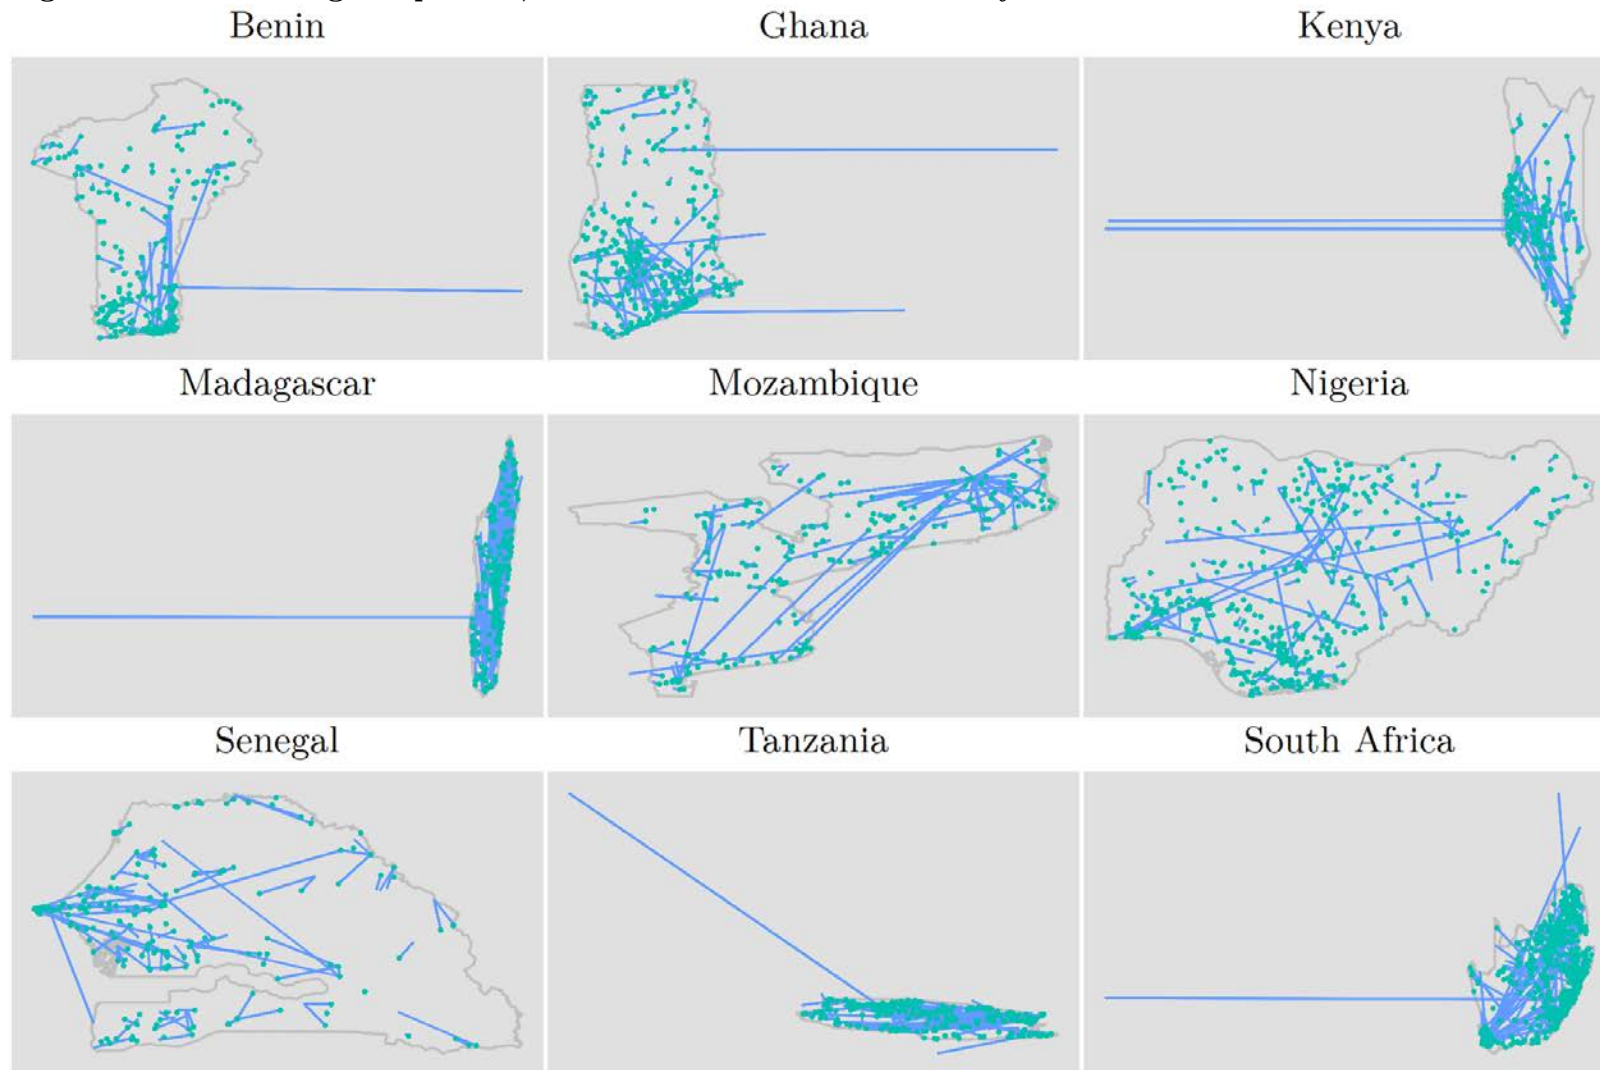
### A Additional figures

**Figure A–1. Street map of part of Cotonou, Benin, with AfTerFibre representation of fiber line**

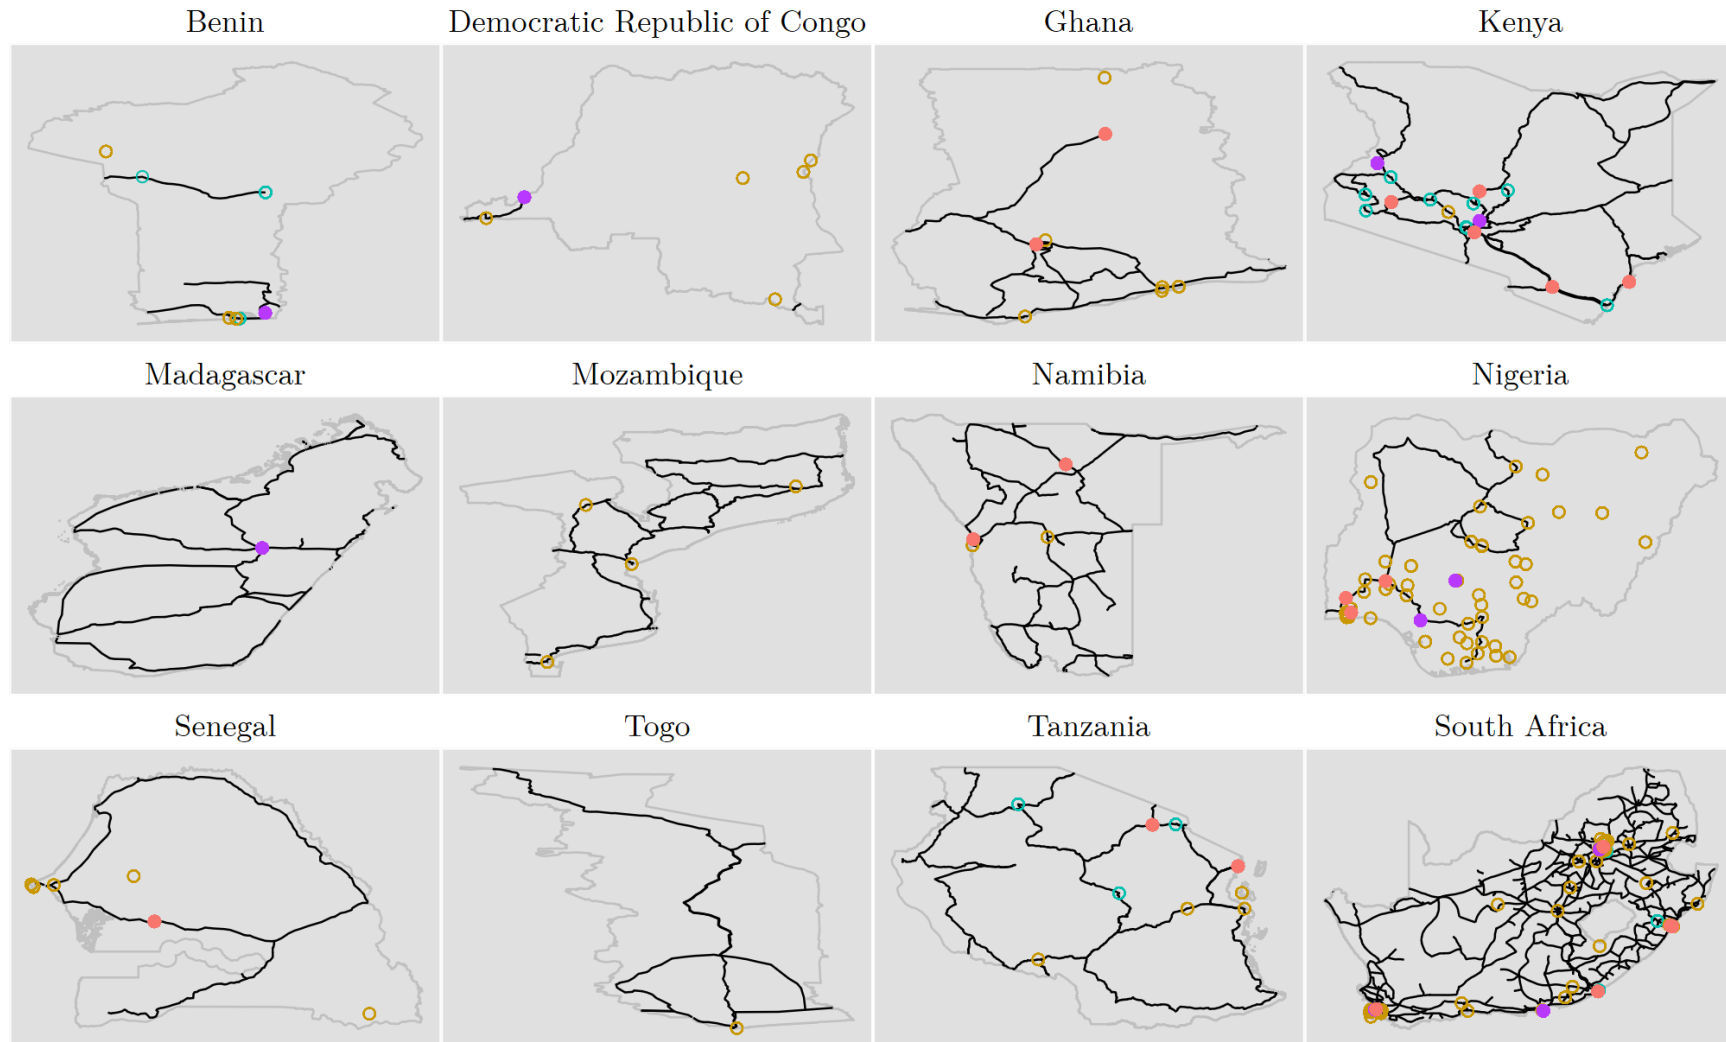


Notes: Map is from OpenStreetMaps ([CC-BY-SA](#)). Superimposed red line shows a fiber conduit as represented in the HP19 fiber network data. Most likely the actual fiber runs along national highway 1 that loops to the south, and only appears to deviate from the highway because its path is based on a digitization of the road route that, at this scale, appears imprecise. The maximum deviation is about  $0.01^\circ$ , twice the  $0.005^\circ$  used in HP19 to define the treatment zone along fiber lines.

Notes: See notes to Figure 3. Only locations successfully re-geocoded through Google Maps are shown, at their newly estimated locations.

**Figure A–2. Geocoding comparison, Afrobarometer household survey locations**

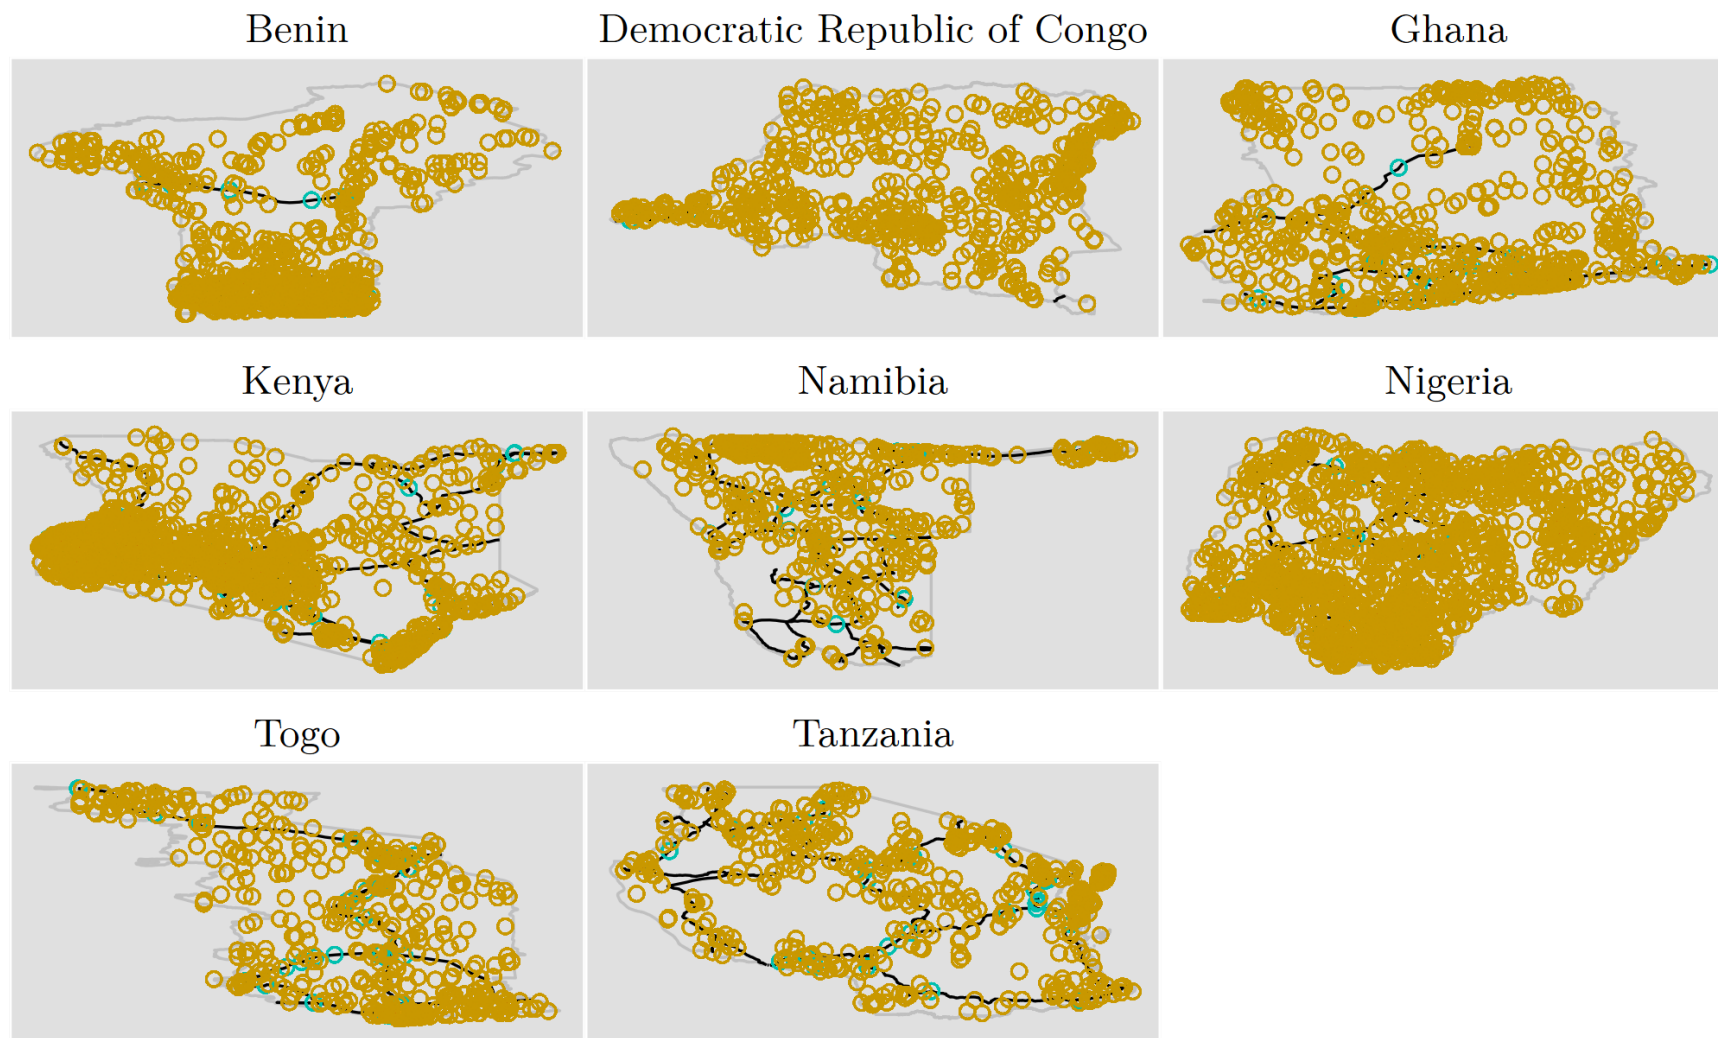
Notes: HP19 and AidData (BenYishay et al. 2017) independently geocode the enumeration areas of the Afrobarometer surveys. These plots compare the two codings by drawing blue lines between the points assigned to each enumeration area and marking the AidData ends with teal dots. Major discrepancies generate long lines with dots at the AidData ends while close matches show just as dots. Scales are chosen separately for each axis and plot. AidData samples limited to locations with precision code=1.

**Figure A–3. Treatment classification comparison, Akamai household survey locations**

Notes: See notes to Figure 3.

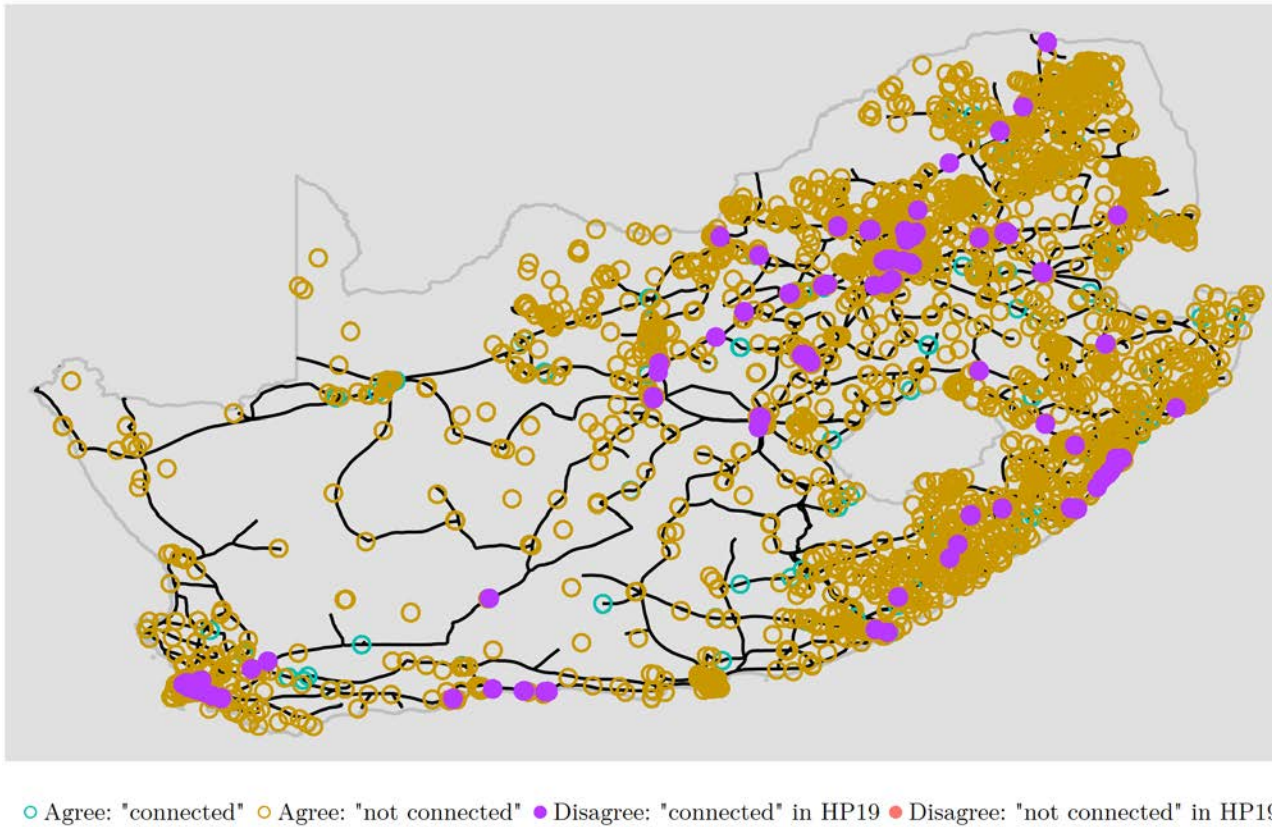


Figure A–4. Treatment classification comparison, DHS household survey locations



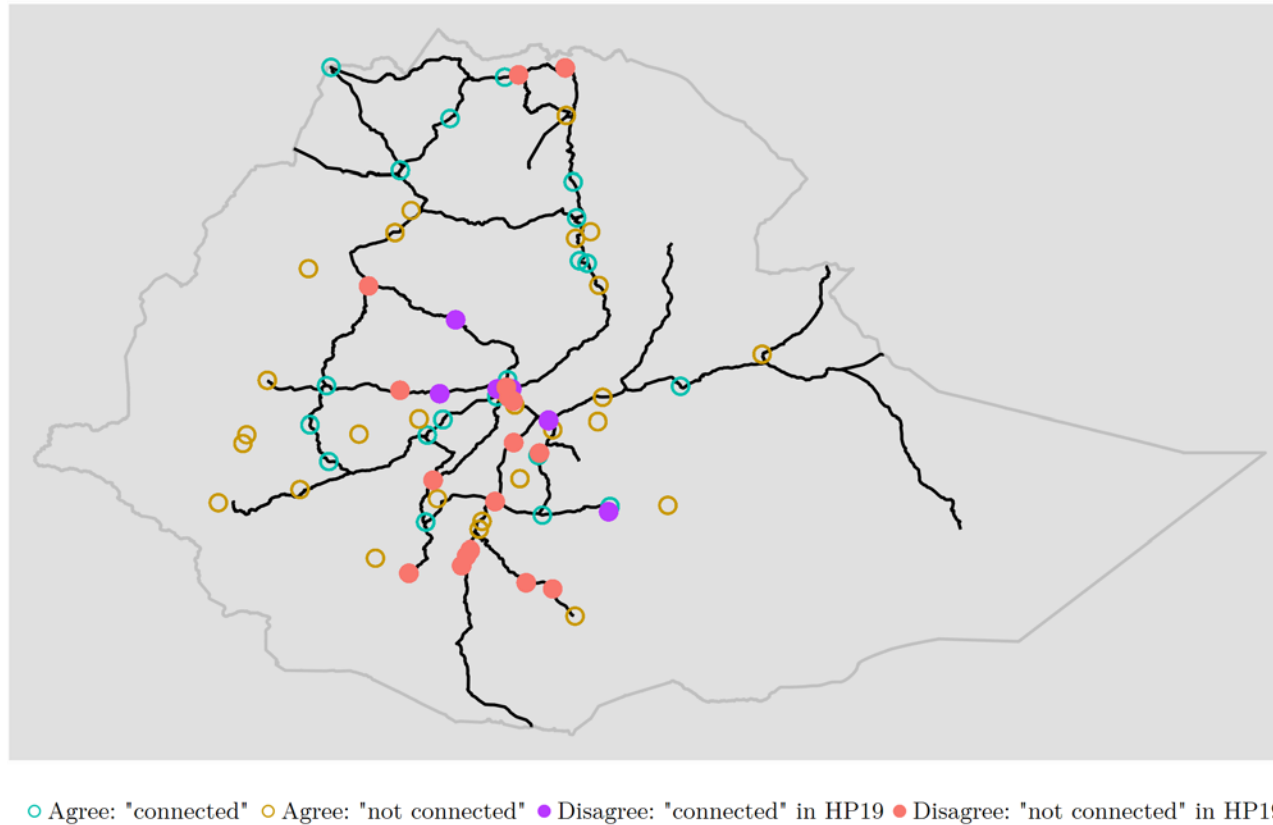
○ Agree: "connected"   ○ Agree: "not connected"   ● Disagree: "connected" in HP19   ● Disagree: "not connected" in HP19

Notes: See notes to Figure 3. There are few disagreements because the HP19 and new classifications use location data provided by the DHS program.

**Figure A–5. Treatment classification comparison, South Africa QLFS enumeration areas**

Notes: See notes to Figure 3.

Figure A–6. Treatment classification comparison, Ethiopia LMMIS firm survey locations



Notes: See notes to Figure 3. The 77 towns and woredas in the data are manually re-geocoded.

## B The econometrics of narrowing the treatment zone

For an exposition on the interaction of measurement error and endogeneity bias, we simplify from difference-in-differences to differences. For an observation  $i$ , conceive of the distance  $d_i^*$  from the backbone as a signed quantity, or coordinate; for example, if a line runs north-south, points to the east could have positive distance and those to the west, negative. Distance is measured with error  $\lambda_i^*$ , yielding the observed distance

$$d_i = d_i^* + \lambda_i^*$$

To dichotomize the notion of treatment, given a treatment zone radius  $r$ , we define true and estimated treatment dummies,

$$\begin{aligned} t_i^* &= \mathbf{1}\{|d_i^*| < r\} \\ t_i &= \mathbf{1}\{|d_i| < r\} \end{aligned}$$

Their difference,

$$\lambda_i = t_i - t_i^*$$

is the error in the measurement of the dichotomous treatment variable.<sup>25</sup>

Assume that in the  $r \rightarrow 0$  limit,

$$\Pr[t_i^* = 1], \Pr[t_i = 1], \text{ and } \Pr[t_i = 1 | t_i^* = 1] \text{ are } O(r) \quad (1)$$

For example, when  $r$  is an extremely narrow 1 meter, the probability that an observation that is actually in the treatment zone will be measured as such ( $\Pr[t_i = 1 | t_i^* = 1]$ ) will be close to zero; and if  $r$  is halved, so will that conditional probability, approximately.

The outcome depends causally on distance according to

$$y_i = f(d_i^*) + \nu_i$$

where  $f$  is a smooth, even function and  $E[\nu_i] = 0$ . Given a value of  $r$ , such as 0.005°, the outcome equation for the associated *dichotomous* treatment indicator

<sup>25</sup> As  $t_i$  and  $t_i^*$  are dummies, the measurement error is non-classical, meaning that  $E[\lambda_i t_i^*] \neq 0$  (Aigner 1973). In particular,  $E[\lambda_i t_i^*] < 0$  because a 0 can only be mis-measured as 1 and vice versa.



is

$$y_i = \beta t_i^* + \epsilon_i \quad (2)$$

where  $\epsilon_i$  is also a mean-zero error and

$$\beta = E[f(d_i^*)|t_i^* = 1] - E[f(d_i^*)|t_i^* = 0]$$

We assume that all variables are iid, which often allows us to drop  $i$  subscripts without introducing ambiguity. But, crucially, we allow the estimated treatment dummy  $t$  to be correlated with the true error  $\epsilon$ , whether because the true treatment dummy  $t^*$  or the measurement error  $\lambda_i$  is also correlated with  $\epsilon$ .

Let  $\hat{\beta}_N$  be the result of ordinary least squares (OLS) regression of  $y$  on  $t$  in  $N$  observations from this process. Substituting the above definitions into the formula for OLS and taking the  $N \rightarrow \infty$  limit while fixing  $r$  gives

$$\text{plim } \hat{\beta}_n = \beta \frac{\text{Cov}[t, t^*]}{\text{Var}[t]} + \frac{\text{Cov}[t, \epsilon]}{\text{Var}[t]} \quad (3)$$

It follows from (1) that

$$\begin{aligned} \text{Cov}[t, t^*] &= E[tt^*] - E[t] E[t^*] \\ &= \Pr[t^* = 1] \Pr[t = 1|t^* = 1] - \Pr[t^* = 1] \Pr[t = 1] \\ &= O(r^2) \\ \text{Var}[t] &= E[t^2] - E[t]^2 \\ &= \Pr[t = 1] - \Pr[t = 1]^2 \\ &= O(r) \end{aligned}$$

Thus the first term on the right of (3)—the exogenous identification component—goes to zero as  $r \rightarrow 0$ . But the second term, the endogenous component, in general does not. In particular, since the latter is the (infeasible) regression of  $\epsilon$  on the dummy  $t$ , we have

$$\lim_{r \rightarrow 0} \frac{\text{Cov}[t, \epsilon]}{\text{Var}[t]} = \lim_{r \rightarrow 0} \{E[\epsilon|d] < r\} - E[\epsilon|d \geq r]\} = E[\epsilon|d = 0] - E[\epsilon] = E[\epsilon|d = 0]$$

If the measured distance  $d$  contains endogenous measurement error, then by definition it contains information about  $\epsilon$ , and the last expression above is  $O(1)$ . So

$$\text{plim } \hat{\beta}_n = E[\epsilon|d = 0]$$

We find that narrowing the treatment zone below the resolution of most the data sources attenuates the exogenous component of  $\hat{\beta}_n$ —the first term on the

right of (3)—toward zero, as in classical measurement error theory. But it does not attenuate the endogenous component, embodied in that second term. While the endogenous component may be second order when the treatment radius  $r$  is large, it can come to dominate when  $r$  is small. As a practical matter, the estimate  $\hat{\beta}_n$  becomes less robust to endogeneity.