Online Appendix (not for publication)

A.I Appendix: Structural Transformation in Ethiopia

Figure A1: Lewis'/Kuznets' Facts on Structural Transformation and Growth

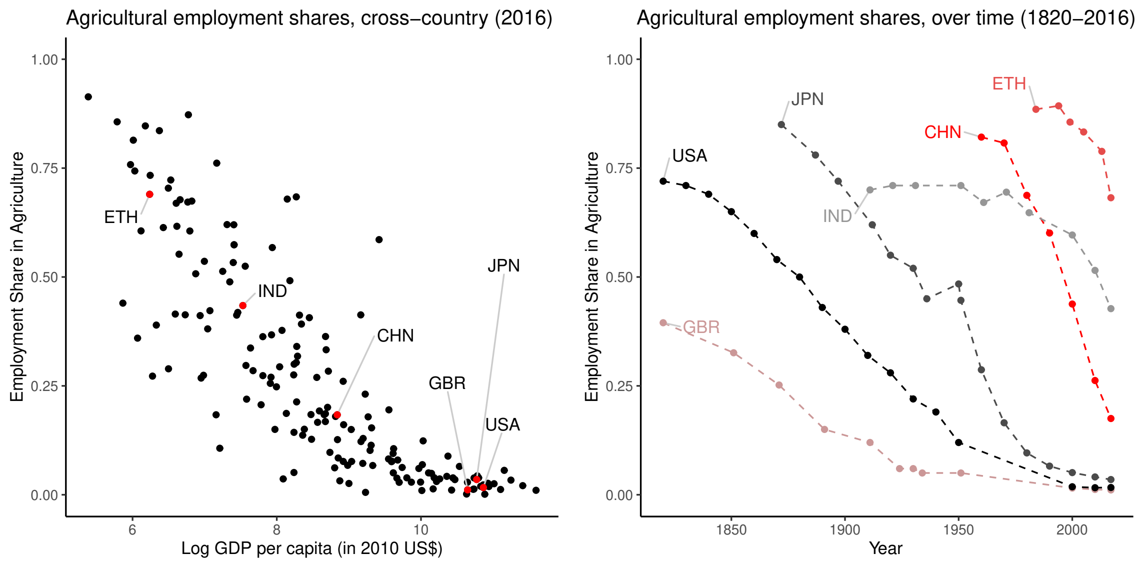

Figure A1 documents Ethiopia's recent structural transformation experience out of agriculture in comparison to cross-sectional evidence across countries as of 2016 (left panel), and in comparison to the time series of selected countries of varying income levels (right panel). Ethiopia currently remains one of the most agrarian economies in the world, although relatively representative of all low-income countries as of 2016. The still high level of employment in the agricultural sector, however, masks dramatic structural transformation over the last two decades, starting from an almost exclusively agrarian society in the late 1980s.

Figure S1 presents additional micro-founded descriptives on macroeconomic structural transformation patterns in Ethiopia that started at least during the mid-1990s, if not earlier. In particular, the share of employment in the agricultural sector declined from 89.3 per cent in 1994 to 56.6 per cent in 2016, despite population growth of approximately two per cent annually, mostly driven by rural, agrarian areas. Starting from very low levels of relative employment, services (manufacturing) employment increased from 7.6 (2.3) per cent in 1994 to 33.5 (9.9) per cent in 2016. Hence, most structural transformation in Ethiopia overall occurred from agriculture to the services sector. However, a comparison of sectoral employment to sectoral value-added trends (see Figure S2) over the same time period highlights a recent uptick in industry value-added between 2011 and 2016, which does not yet appear to result in markedly higher relative industry sector employment.

Especially if structural transformation is of a low-level nature, i.e. out of agriculture into mostly small-scale, informal retail services (see Section 4), positive income and welfare effects of such sectoral shifts are not obvious, neither at the individual level, nor in the aggregate. However, as shown in Figure S3, my study period displays an almost exploding time series of GDP per capita and a dramatic reduction in headcount poverty, using either national or international measures. Hence, even from a purely descriptive perspective, a study of large-scale infrastructure investments in relation to structural transformation seems warranted.

A.II Appendix: Sample Coverage for Reduced-Form

Figure A2: Sampling of Districts, Combined NLFS and DHS (1999-2016)

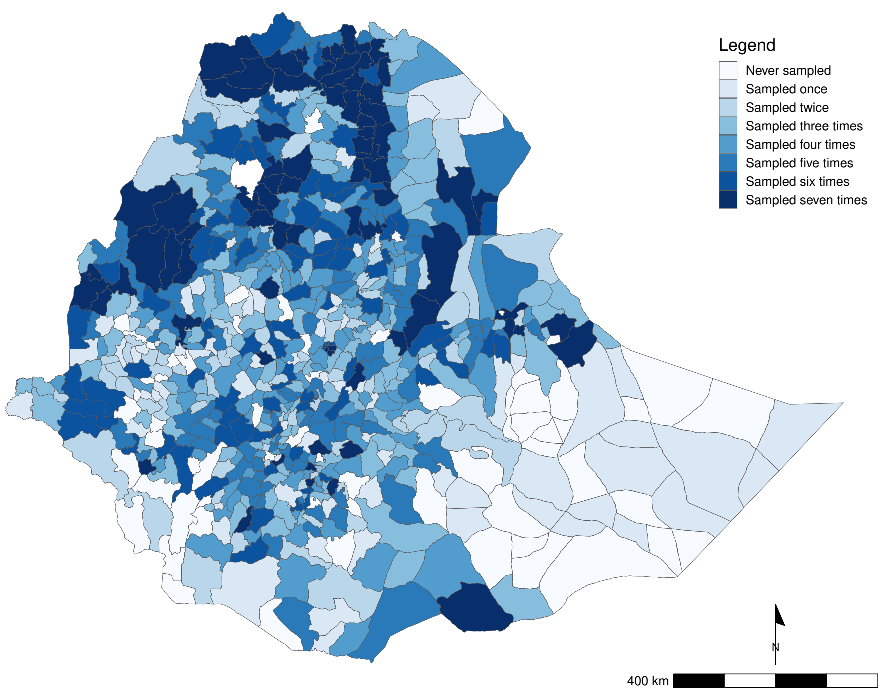

Figure A2 confirms graphically which districts where covered how many times by the two largest household surveys, the high-quality and internationally standardised Demographic & Health Survey (DHS) and the larger Ethiopian National Labour Force Survey (NLFS), all conducted by Ethiopia's Central Statistical Agency. In total, four rounds of survey data of the Demographic & Health Survey (DHS) are available [2000, 2005, 2011, 2016]. These repeated cross-sections of household-level (and invidiual-level) data are complemented by three rounds of the Ethiopian National Labour Force Survey (NLFS) [1999, 2005, 2013], which yields a high combined coverage of my study period of interest from the late 1990s to the recent past. As Figure A2 shows, districts can be sampled up to seven times in this combined sample used for the main reduced-form analysis and results in Section 4. Three areas of Ethiopia remain uncovered or severely undercovered: the Somali region in the Southeast, a mostly desert area scarcely inhabited by pastoralists that are often excluded from CSA surveys; the dense forest areas in the Southwest, which are mostly void of population; and the Danakil depression area and desert in the Northeast, at the border with Eritrea, for similar reasons. The DHS and NLFS combined provide the most frequent and largest spatial coverage for occupational choice data in the case of Ethiopia, and neatly cover the crucial two decades from the late 1990s to the late 2010s when roads and electricity networks expanded fastest in Ethiopia. An additional survey instrument covering household consumption, the Household Consumption and Expenditure Survey (HCES), is also available for four rounds [1999, 2005, 2011, 2016]. This broad spatial and temporal coverage across survey instruments presents unusually rich secondary data in a low income context.[87]

A.III Appendix: Infrastructure Expansion in Ethiopia

Figure A3 documents how the all-weather road network expanded roughly fourfold between the late 1990s and 2016, from approx. 16,000km to approx. 70,000km. The focus on all-weather roads, i.e. roads with either asphalt, bitumen or gravel surface, arises from the understanding that trade and market access rely on a given location's year-round accessibility (ideally by lorry). For 1999, no distinction between

[87] Ethiopia conducted population censi in 1984, 1994 and 2007, with the 2017 census experiencing ongoing delays. The mid-study 2007 census round saw the occupational choice question dropped.

Figure A3: Large-scale Road Network Expansion (1999-2016)

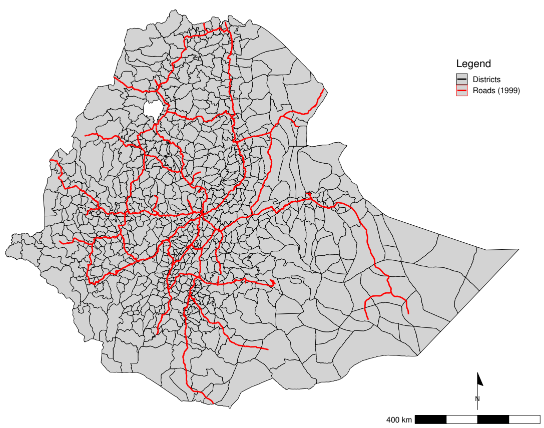

Figure A4: Large-scale Electricity Network Expansion (1990-2016)

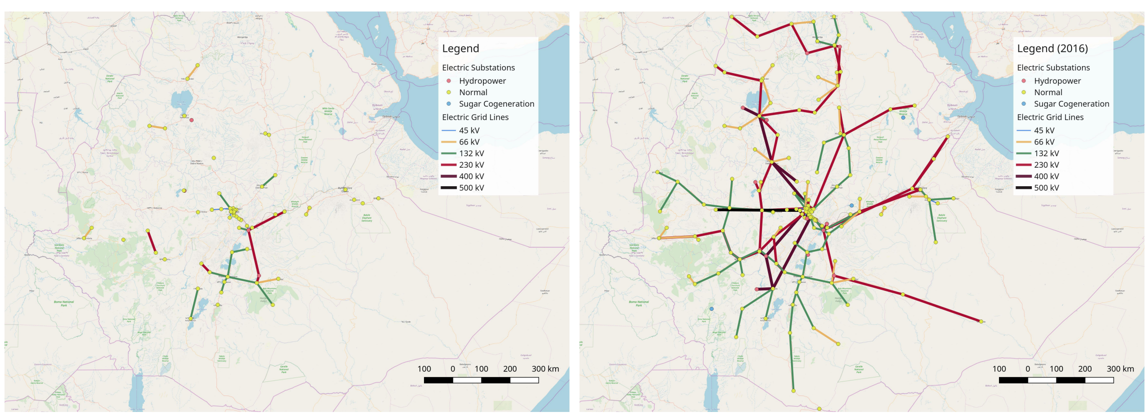

asphalt, bitumen and gravel was available, whereas for 2016 asphalt/bitumen roads are highlighted in red, whereas gravel roads are highlighted in blue.

Figure A4 documents how over the same period the electricity network doubled from 95 to 191 major electric substations. Substations, highlighted as yellow points, are crucial for electrification since they step down high voltages used in long-distance transmission to low voltages used for local distribution networks. High voltage transmission lines efficiently conduct electricity over long distances between major sources of generation (such as hydropower dams) and demand concentrations (such as cities). Local low voltage distribution networks supply individual firms and households.

The almost complete lack of direct infrastructure substitutes in Ethiopia implies that the all-weather road and electricity network expansions I track capture genuine extensive margin effects of access to infrastructure. In particular, Ethiopia is a land-locked country without major navigable rivers or canals. During my study period, the single existing railway line (to neighbouring Djibouti and its port) was still out of order.[88] Another new railway project only began construction in 2015.[89]

With respect to access to energy and substitutes for grid electricity, only a handful of isolated diesel generators originating from the 1960s operated in selected major cities. All of these major cities were grid-electrified before my study period and, thus, do not feature as compliers in the empirical strategy below. Self-generated energy from off-grid solar home systems generated approximately one megawatt of capacity midway through my sample, compared to total installed grid capacity in 2018 of 4,256 MW.[90] Thus, due to low penetration and low voltage in Ethiopia during my study period, off-grid solar cannot be regarded as a feasible substitute to grid electricity access. Other off-grid generation alternatives (such as mini-hydropower systems) are not known to have been present beyond isolated cases.

One potential concern with tracking infrastructure expansion in a country with a centrally located capital city (which also happens to be the country's largest, as well as its administrative, commercial and industry hub) may be the expectation that economic activity, in line with population density, decreases radially from the centre. In this case, any reasonable least-cost infrastructure expansion should also follow a radial process outwards from the centre, allowing location and timing of expansion investments to be expressed as a simple function of distance to the centre. Fortunately, this hypothesis can be rejected in Ethiopia: Figure S9 shows that population density in Ethiopia is spread out irregularly, and also does not interact in a straightforward manner with either elevation (see Figure S10, left panel) or terrain (ibid., right panel). Large parts of the Ethiopian population live in highly rugged, elevated and remote locations, which do not align with radial distance to the economic centre.

[88] A recently completed, newly built replacement railway to Djibouti was inaugurated in October 2016. Due to equipment failures, however, commercial operations only started in January 2018.

[89] cf. International Rail Journal's news coverage in February 2015: https://www.railjournal.com/index.php/africa/work-starts-on-delayed-ethiopian-project.html

[90] An additional 25,000 solar home panels (à 5-10 W) were purchased by the Ethiopian government for decentralised installation by 2013, cf. https://allafrica.com/stories/201308070099.html

A.IV Appendix: Background on Data Collection

Resulting from a close collaboration with Ethiopian Electric Power (EEP), the state utility charged with electricity generation and transmission, I obtained confidential information on the exact location, capacity, equipment and commissioning times of each of the electric grid's substations.[91] These records cover a total of 191 substations which came online before 2018, with the first isolated substations constructed in 1959. To reconstruct the expansion of the interconnected system ("the grid"), I also obtained each transmission line's location, connecting nodes, voltage and commissioning times, as well as further information regarding recent upgrades into stability-enhancing equipment (e.g. reactors, capacitators) along each line.[92]

Finally, I also collected locations, capacity, operational status and commissioning time information on all power plants to track generation. The Ethiopian electricity supply is mostly provided by hydropower from nine major dams, as well as at least three wind farms, one geothermal power plant and by-generation from at least three sugar refineries. Dam openings since 2016 are currently ignored in my analysis due to the lack of outcome variables spanning this very recent past.

Data on all-weather roads were obtained from the Ethiopian Roads Authority (ERA). In particular, I employ several historical maps from various, partially undisclosed records. For the years 2006, 2012 and 2016, I have obtained GIS data, which rely at least partially on actual road surveys: in particular, the final cross-section from 2016 relies on a months-long on-the-ground data collection by ERA that verifiably mapped every road in the country, recording surface type, quality, width, current state and GPS markers at regular intervals. Earlier maps were supposedly based on partial road surveys or construction project records. However, given the lack of central records documenting project-level road construction, I cannot verify this claim.

In addition, I use various other sources for cross-validation and better visibility on the pre-sample road network: I use GIS data from OpenStreetMap for the year 2014 to cross-verify the earlier and later ERA records. I also use manually digitised

[91] Formerly a single state utility, Ethiopian Electric Power Corporation (EEPCo), EEPCo was broken up into two separate entities in 2013: a generation and transmission utility, Ethiopian Electric Power (EEP), and a distribution utility, Ethiopian Electric Utility (EEU).

[92] To inform cost-benefit calculations, I also collected construction cost estimates from the engineering team for unit costs of transmission lines and past records of selected actual project expenditures.

historical CIA maps from 1969, 1972, 1976, 1990 to obtain the pre-sample period. The CIA's 1999 map is used as the first cross-section in the sample and the 2009 map for cross-validation of ERA records. Furthermore, I also make use of a biennial, district-level road density dataset (1996-2012) kindly provided by Shiferaw et al. (2015) for robustness. Changes in district road density correlate highly with map-derived district-level all-weather road access I employ in the main analysis.

With respect to the geo-identification of enumeration areas (and, thus, households) required for the outcome variables, two qualifications are due: first, the enumeration area locations of NLFS EAs are provided in codified form, which is at times only imperfectly geographically traceable. Missing codebooks at the Ethiopian Central Statistical Agency in combination with missing old maps make cross-referencing of old codebooks to old district and enumeration area delineations for some cases close to impossible. Second, even the DHS-provided GPS coordinates for EAs locations are not perfectly reliable due to the common random displacement applied to GPS coordinates prior to publication. To ensure survey respondents' anonymity, DHS EA coordinates of rural (urban) EAs are randomly displaced within a 0-10km (0-5km) radius.[93] Therefore, although I have exact geo-identified information on infrastructure placement, the finest spatial resolution the analysis can support is at the district-level due to geo-identification constraints arising from the available outcome variable data.

Finally, for reduced-form analyses of household expenditure responses to infrastructure investments, I obtained several rounds of the CSA's Household Consumption and Expenditure Surveys (HCES) were obtained – in particular the 1999, 2005, 2011 and 2016 rounds. Of these, the 1999 and 2016 rounds can be geo-located. For 1999 (2016), 1,264 (2,106) enumeration areas were sampled and 17,332 (30,229) households surveyed, while, as is common for most household surveys in Ethiopia, the non-sedentary regions (Afar, Somali) were excluded. From these two geo-located survey rounds I obtain information on household expenditure per capita, household size and household demographics. Aggregated from households to enumeration areas, repeated draws of enumeration areas from the same district over time allow creation of a pseudo-panel at the district-year-level similar to the main DHS/NLFS pseudo-panel.

[93]In theory, displacement of EAs should never cross zone borders (the second highest administrative level), nor country borders, although they may cross district borders (the third highest administrative level). In practice, a handful of displacement errors are corrected manually.

A.V Appendix: Endogenous Infrastructure & IV Creation

Figure A5: Endogenous Allocation of Treatments across Districts (as of 2016)

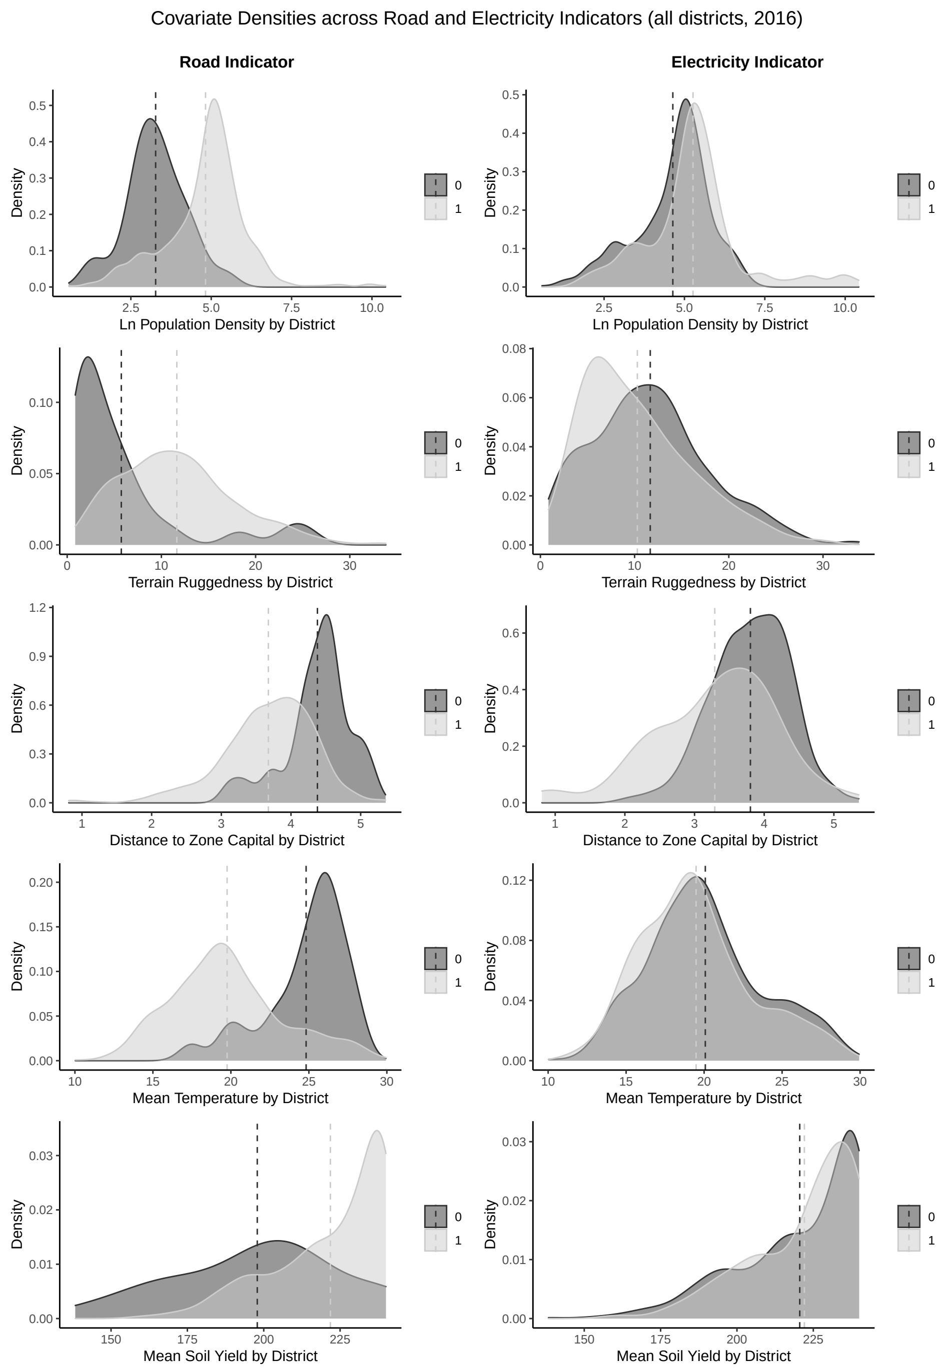

In the following, I discuss three identification challenges that seem particularly relevant in the Ethiopian context: potential spatial targeting of investments across districts, temporal prioritisation of investments across districts, and natural sequencing of different investments arising from interdependence between road and electricity network investments. I then briefly discuss more classic sources of endogeneity and the degree to which they may apply in Ethiopia during the study period.

A classic endogeneity concern in the allocation of infrastructure investments is that policymakers make active allocation decisions in which locations should obtain access and which not. Figure A5 confirms endogenous allocation of infrastructure in Ethiopia based along various observables. This is not surprising: the high cost of transmission infrastructure makes spatial targeting of locations to receive access is an obvious feature of electrification. In particular, existing engineering guidelines highlight the primary cost driver to be minimised as the length of transmission lines. Engineers involved in the Ethiopian grid expansion have also reported privately that in order to minimise cost, during the early years of electrification only politically demanded locations would obtain a transmission line connection (and substation), apart from locations that accidentally happened to lie on a relatively straight line between supply and demand. Despite recent electrification advances, many rural and remote locations will remain without electricity access for the foreseeable future.

Another cause for endogeneity in the allocation of infrastructure investments across locations could be temporal prioritisation according to unobservables if all locations eventually obtain infrastructure access. In the particular case of roads in Ethiopia, for example, the government formulated an explicit policy to connect all of the (at the time) 689 district capitals with an all-weather road by 2020 – an objective that was successfully achieved by 2016 already. Hence, in my analysis of road network investments, a key endogeneity concern is the timing of a district's connection (in contrast to the above issue of endogenous district selection into treatment, since all districts obtain road access treatment eventually). Figure S4 confirms that more densely populated districts receive an all-weather road earlier than more rural, sparsely populated districts, pointing towards endogenous connection timing.

Finally, at the level of transmission lines and substations, these major infrastructure items also crucially rely on at least some means of transport to be available for construction. Hence, sequencing of infrastructure investments appears natural in the context of transport and electrification. Translated to the Ethiopian context of big push infrastructure investments, a possible endogeneity concern arises from the fact that electrification only reaches previously road-connected places.

Further identification challenges arise from other sources of omitted variable bias that affect both the infrastructure expansions and structural transformation (such as natural resource windfalls, global economic cycles, capital flows, donor funding, etc.), which are entirely plausible in the Ethiopian context. Likewise, reverse causality in the form of sectoral shifts causing greater demand for infrastructure investments should also not be ruled out *ex ante*. Finally, measurement error in the right-hand side variable may lead to attenuation bias, for example due to inaccurate timing information of electric grid expansion or imprecise historic road maps.

To address the above endogeneity concerns, I present two instrumental variables. Regarding electrification, the instrumental variable (IV) is founded on the fact that electricity supply must be connected to demand, or in engineering terms: to the load centre. Translated to the Ethiopian context, electricity generation originates to 98 per cent of total installed capacity from hydropower dams in the Ethiopian highlands. The largest load centre, however, is Addis Abeba, which also hosts the load dispatch center of the interconnected system which manages supply allocations.

Therefore, I develop an IV which yields a hypothetical electrification status and timing for each location based on that location's proximity to a straight line corridor from a newly opened hydropower dam in mostly remote parts of Ethiopia to Addis Abeba. From the year of dam opening onwards, all districts lying along this straight line connecting the dam to Addis will be considered hypothetically electrified.

With respect to such an IV's identifying assumptions, the validity assumption reads that the hypothetical electrification status of districts along a straight line from a new dam to Addis does display a statistically significant relationship with these districts' actual electrification status and year of electrification. I draw straight line connection corridors of 25km diameter for nine dams and two large-scale wind farms.

The random assignment assumption of the IV would imply that a given district's exposure to a straight line corridor was spatially and temporally as good as randomly assigned. In other words, locations that lie between both of the straight line endpoints, which would usually span several hundred kilometres, are not systematically different from nearby locations off the straight line corridor. Likewise, the timing of the high-voltage line coming online due to the opening of the hydropower dam should also be exogenous. Given frequent multi-year delays in these large dam construction projects, the assumption of exogenous final commissioning time appears to hold in Ethiopia.[94]

Finally, the exclusion restriction requires that the straight line corridor from dam to Addis does not affect structural transformation in the years and locations exposed to the hypothetical transmission line, other than through actual electrification.

In sum, my electrification IV represents a classic 'inconsequential units' IV (cf. Redding and Turner (2014)) brought to the electrification context.[95] Figure A6 provides a graphical representation of how proximity to straight line corridors (and their opening years) generate spatial and temporal variation in districts' predicted electrification.

Regarding my instrumental variable for the timing of a district's road connection, I construct a hypothetical road expansion based on distance to straight line arteries drawn up by 1930s Italian colonial plans for road construction: In order to conquer Eritrea, Ethiopia and Somalia, as well as to effectively occupy their territory, the Italian invaders initiated a large-scale road construction effort starting in 1936. Either lacking information about the local geography and terrain or actively ignoring it, Benito Mussolini himself appears to have designed at least five major road arteries to connect the capitals of former ancient kingdoms to each other and to major ports, enabling the Italian colonial forces in theory to project power into conquered territory.

In particular, straight line axes were drawn to connect Addis Abeba, the capital of the defeated Ethiopian Empire, to both Asmara (then capital of Italian Eritrea) and Mogadishu (then capital of Italian Somaliland). In addition, the ancient kingdom capitals (and centres of regional power) of Gonder (Begemder Kingdom), Dessie (Wollo province), Nekempte (Welega province), Jimma (Kaffa Kingdom), Yirga Alem (Sidamo Kingdom) and Harar/Jijiga (Emirat of Harar/Hararghe province), as well as the Red Sea port at Assab were to be connected either directly to one of the major capitals or on the way. The resulting straight line arteries are depicted in Figure A7.

[94] In private conversations with current and former EEPCo and EEP senior engineers in charge of grid planning and expansion, i.e. local experts with decades of relevant sector-specific experience, providing accurate predictions of dam construction delays was described as non-trivial.

[95] Examples of similar instrumental variables include Michaels (2008) and Kassem (2018), who also use exposure to artificial lines to instrument for infrastructure expansions.

Starting in 1936, actual Italian road construction followed to a surprising degree Mussolini's grand design of unrealistically straight road arteries, irrespective of the adverse terrain covered. Before their defeat at the hands of British and allied forces in the Horn of Africa in 1941, Italian colonial authorities managed to construct at least 4,000 kilometres of paved and 4,400 kilometres of unpaved roads. Figure S11 provides a historic picture of the construction efforts during the late 1930s.

On the territory of today's Ethiopia, approximately 3,378 kilometres of paved 'highways' were constructed, of which at least 1,970 kilometres were finished including state-of-the-art asphalt surfacing. Importantly, a lasting feature for future Ethiopian road construction were the 4,448 small and 128 large bridges finished by the Italian colonial authorities, artefacts necessitated by the idiosyncratic routing through the Ethiopian Highlands mass and multiple mountain ranges.[96]

For the purposes of the roads IV, I exploit the fact that Ethiopian road construction in the 1990s started reconstruction of roads from the former Italian colonial trunk network and subsequently, during my study period from 1999 to 2016 fanned out road access almost orthogonally to nearby cities, towns and settlements, closely following geographic features (i.e. mostly valleys and ridges). Figures S12 and S13 provide two exemplary cases of how Ethiopian road construction connected nearby locations and districts almost orthogonally starting from the (reconstructed) Italian colonial roads.

I therefore construct a roads IV in the following way. Starting from the seven straight-line arteries designed by Mussolini (and depicted in Figure A7), I calculate orthogonal, shortest distances to every district capital, as the crow flies.[97] One should also note that this distance is calculated from the plausibly exogenous straight lines schematically drawn by the Italians, not from the actual roads that were subsequently constructed (and re-constructed) under these designs. Since road construction in today's Federal Democratic Republic is a politically regional matter, I run the following algorithm separately and simultaneously for each of the eleven regions of Ethiopia.[98]

[96] Apart from the vast Ethiopian Highlands ('roof of Africa'), of the remaining eight major mountain ranges, four were crossed: the Ahmar, Entoto, Semien Mountains, and the Mount Afdem range.

[97] Districts containing arteries are considered already 'treated' with road access by the roads IV.

[98] For the city-regions of Addis Abeba and Harar, in which districts always had access to at least one all-weather road in 1999, the instrument predicts road connections in 1999 already. For the city-region of Dire Dawa, the instrument predicts both of the city's districts to be connected by 2000. The algorithm procedure to generate temporal variation therefore only generates meaningful

One relevant peculiarity of the algorithm is that road distances to connect a given district are never updated: the calculated distance is always taken as the distance to the nearest Italian artery, which does not vary over time. This deliberate choice against a continuous-updating algorithm, that would calculate the shortest distance to either the Italian artery or the nearest connected district capital, arises from a potential threat to the exclusion restriction, where short district connections from district capital to district capital pick up agglomerations of population (and thus smaller district sizes). Once the district closest to the artery of such an agglomeration would be connected, the succeeding districts would be connected relatively sooner compared to an algorithm without continuous distance-updating. Therefore, to guard against this potential violation of the exclusion restriction, I do not update distances and always have the algorithm build (relatively less realistic) connections to the closest Italian artery, irrespective of any districts already connected in between.

In sum, my Italian artery roads IV provides temporal and spatial variation in district road access derived from a plausibly exogenous source, namely orthogonal straight line distance to Italian straight line arteries. Figure A7 depicts the resulting district-level variation in the predicted arrival year of an all-weather road.[99]

For robustness, I also generate an alternative instrumental variable for a district's road connection derived from Boruvka's instead of Kruskal's algorithm, following the explicit policy objective to connect all district capitals by the end of the sample period in 2016. The algorithms thus provide spatial variation in terms of how each district will get connected to the network (which vary slightly across Kruskal's and Boruvka's algorithms with respect to their order), but do not yet provide any temporal variation in districts' connection timing. Therefore, I additionally employ the above budget split rule to the output of Boruvka's minimum spanning tree, such that only a certain amount of new all-weather road mileage can be built per year, following the order dictated by the algorithm. I obtain an alternative hypothetical road network that features both spatial and temporal variation. Results are qualitatively unchanged and very similar in quantitative magnitude across specifications.

variation for the remaining eight major regions of Ethiopia: Afar, Amhara, Beneshangul Gumuz, Gambela, Oromia, SNNPR, Somali and Tigray. Given the status of the three city-regions as 'always takers' in the frequentist sense, this prediction is expected and resonates well with empirical reality.

[99] Accordingly, the instrument takes a value of one from the district-year in which a given district gets connected (as determined by the regional budget split rule) onwards.

Figure A6: Electrification IV Corridors and Times, Connecting Dams with Addis

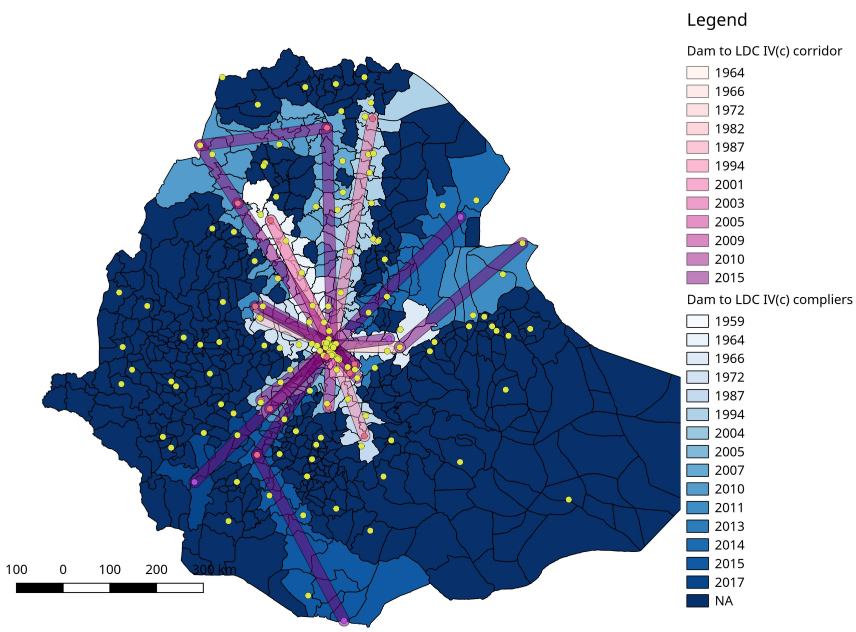

Figure A7: Road IV (Italian-Kruskal) District Connection Year to All-weather Road

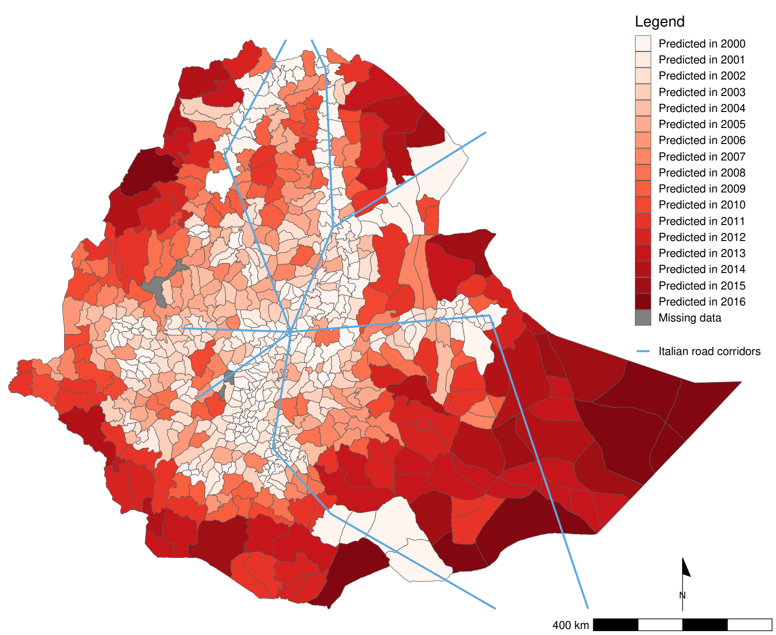

Figure A8: Quasi-Random Assignment of Electricity IV: Original vs 'Next 25km Buffer'

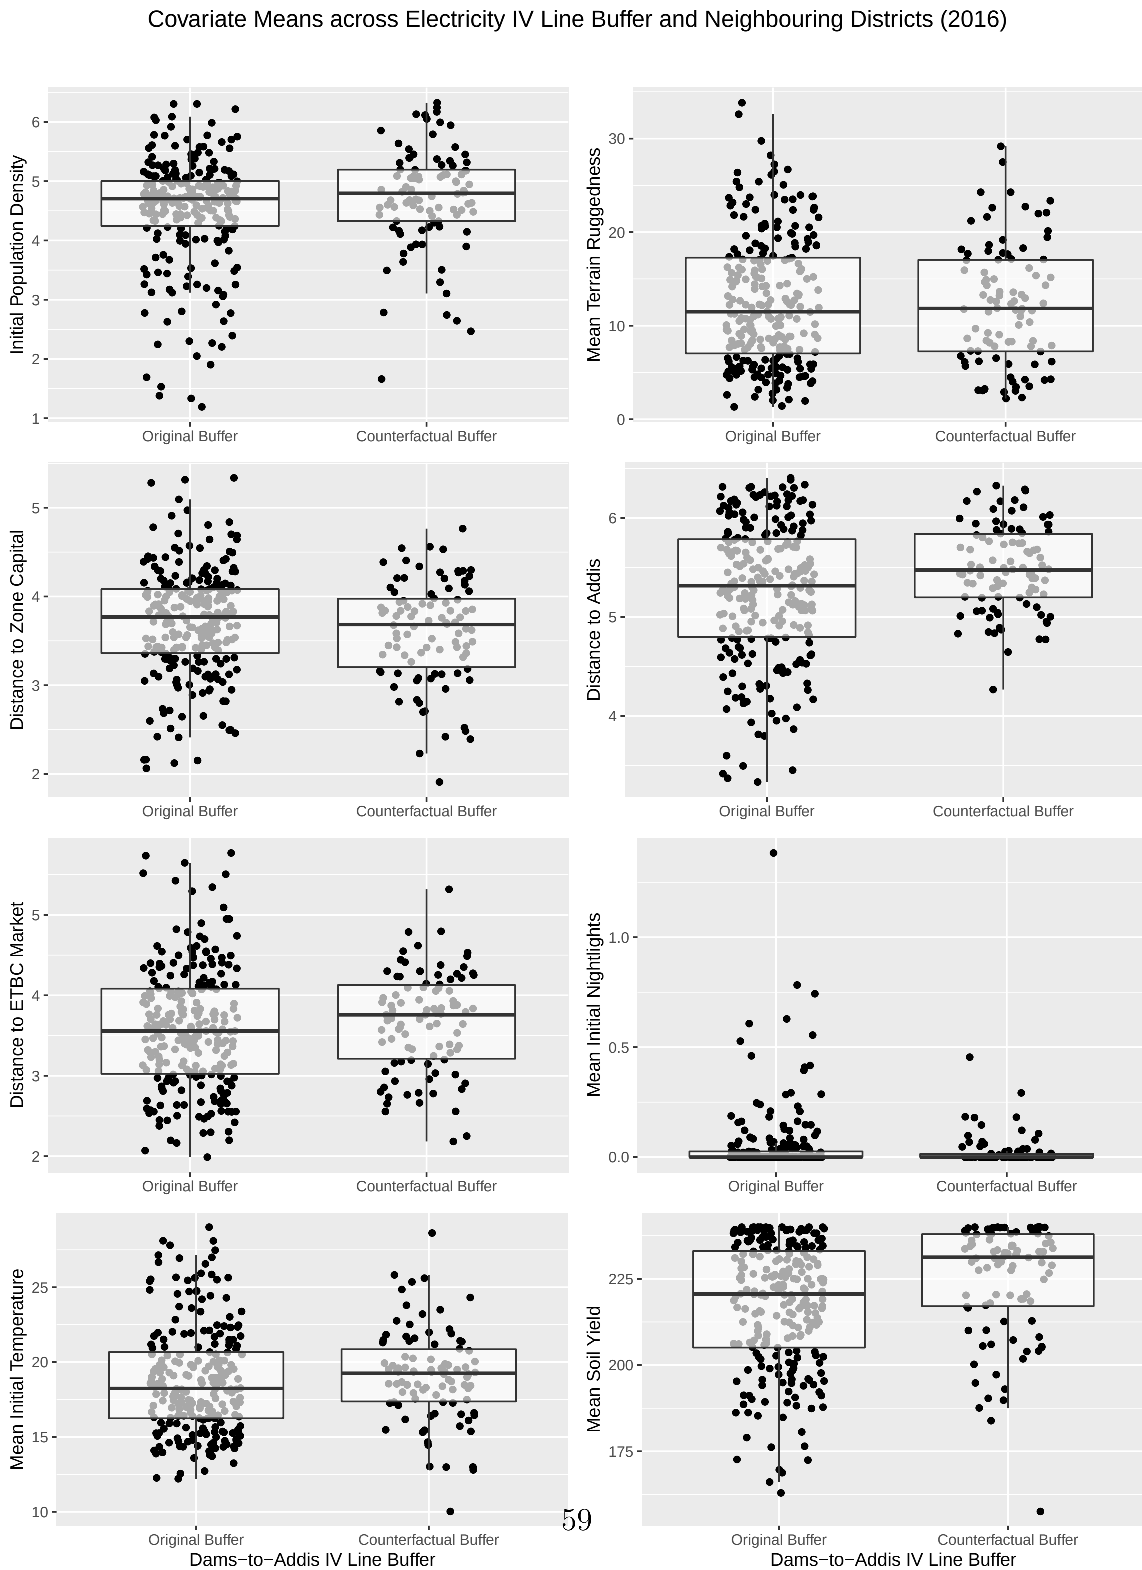

Figure A9: Quasi-Random Assignment of Roads IV: Predicted Connection Just Before vs Just After Survey Years

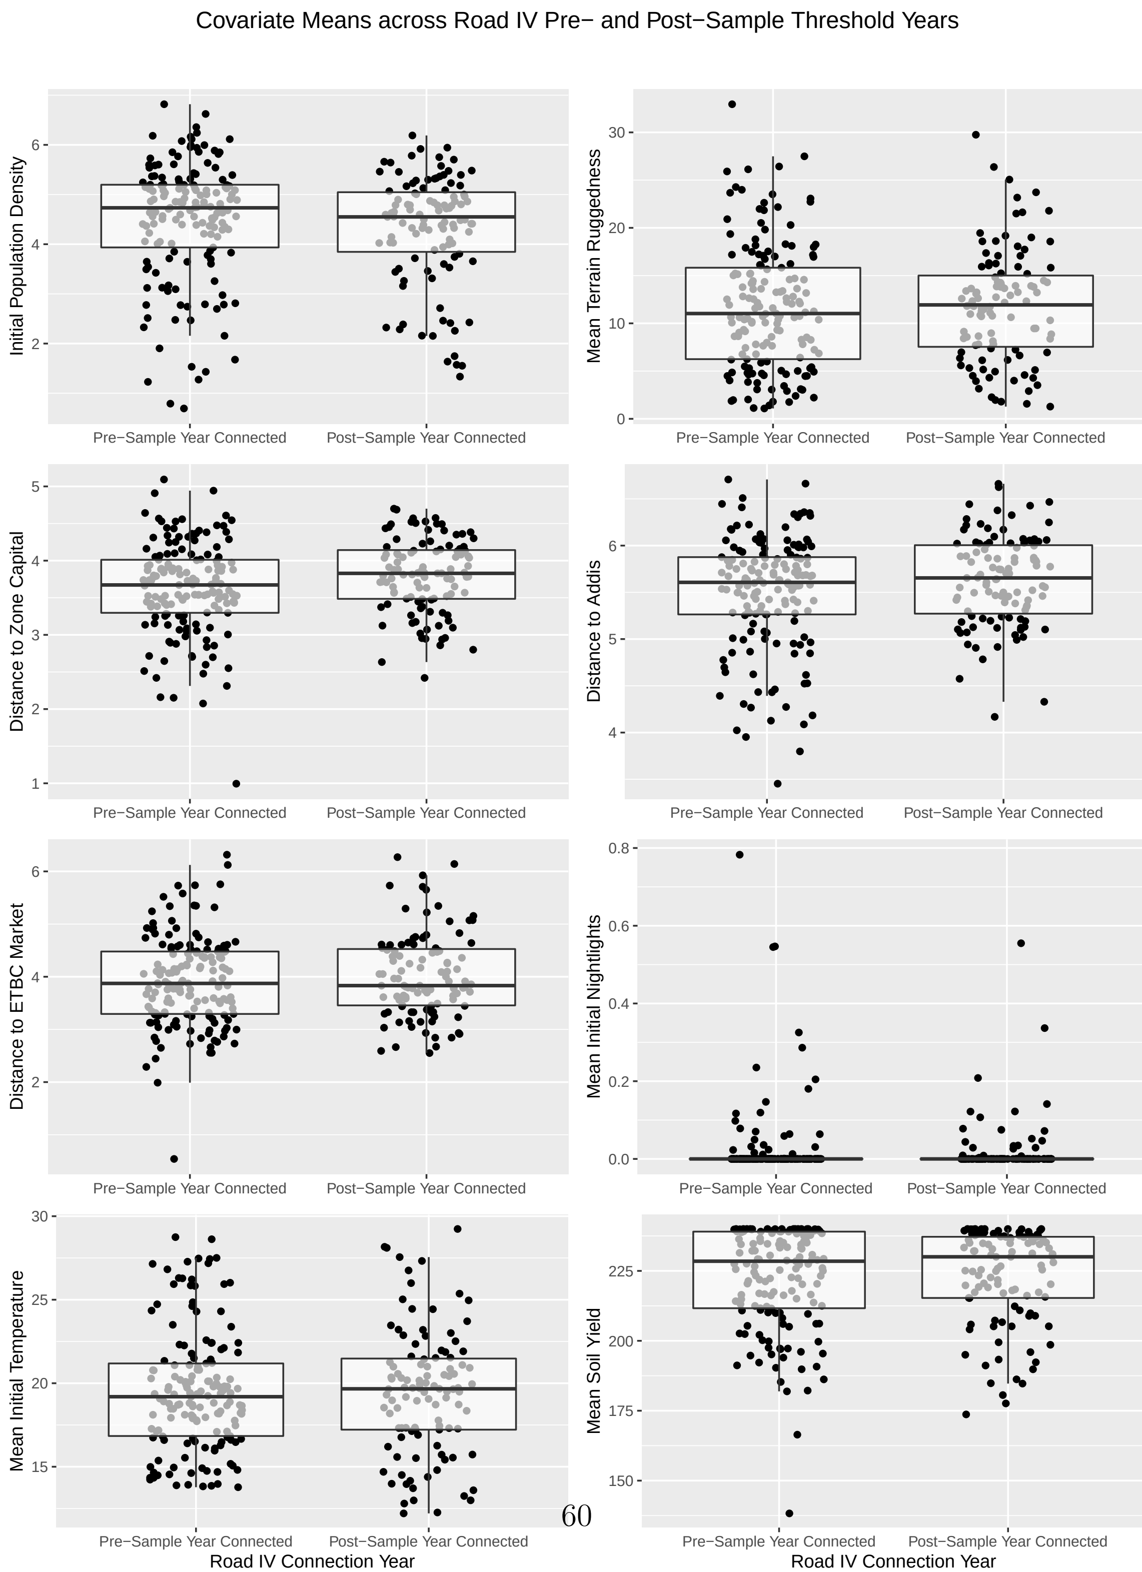

A.VI Appendix: Additional Reduced-Form Results

Figure A10: Second Stage: Occupat. Change (ISCO) on Roads & Elec., full sample

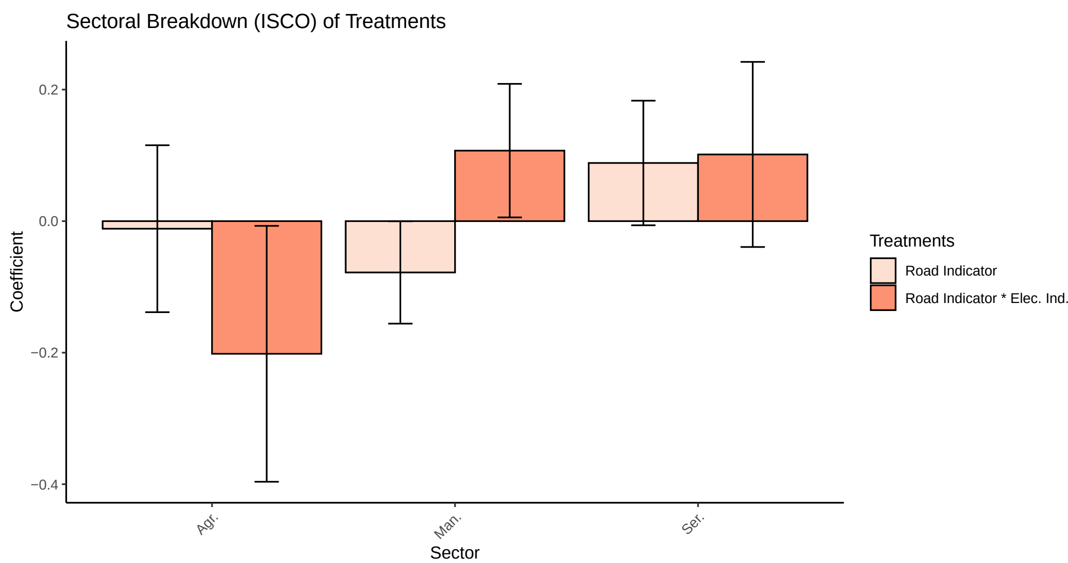

Figure A11: Second Stage: Occupat. Change (ISCO-1d) on Roads & Elec., full sample

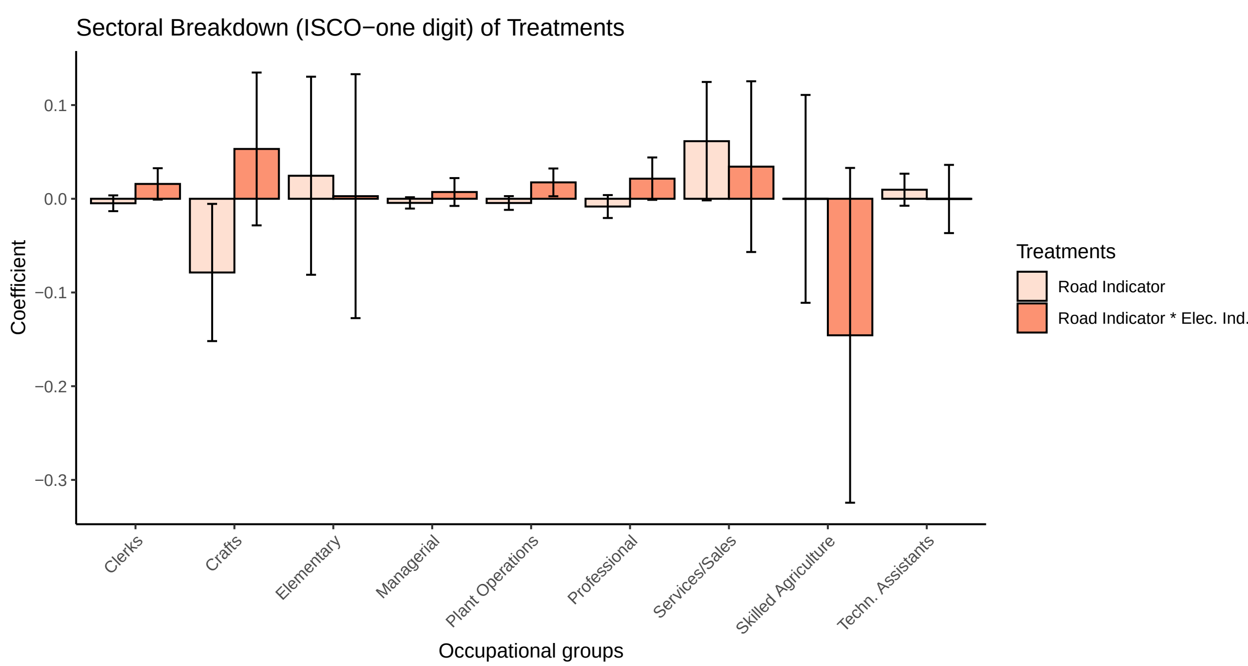

Table A1: Second Stage: Occupat. Change on Roads and Elec., full sample (1999-2016)

	Agriculture	Services	Manufacturing
	(1)	(2)	(3)
Ln(Market Access)	-18.340	74.843	-57.117
	(48.702)	(47.425)	(38.916)
Ln(Market Access)*Elec Ind.	-0.005^{**}	0.003^{*}	0.002^{*}
	(0.002)	(0.002)	(0.001)
Sample	*full*	*full*	*full*
Year FE	✓	✓	✓
Region-Year Trend	✓	✓	✓
Controls	✓	✓	✓
Windmeijer cond. F.	5.852	15.932	
p-val $\beta_1 + \beta_2 = 0$	0.707	0.115	0.142
p-val $\beta_1 - \beta_2 = 0$	0.707	0.115	0.142
R^2	0.293	-0.027	-0.249
Adjusted R^2	0.263	-0.070	-0.301
Observations	2007	2007	2007

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. All standard errors clustered at district-level. Sample includes all districts sampled at least twice in either NLFS (1999, 2005, 2013; 1163 district-years) or DHS (2000, 2005, 2011, 2016; 999 district-years), yielding a combined total of 2007 district-years. All districts in Addis Abeba region are always excluded. Individual answers to survey question about 'current occupation' grouped into sectors according to ISCO major groups. 'Agriculture' denotes aggregated share of district-year survey sample of working age that respond to work in an agricultural occupation. Omitting results for outcome variable 'Currently Not Working', which denotes share of survey sample of working age not currently working. Log market access and log market access–electricity indicator interaction denote predicted values from the two first stage specifications. Region-year-specific time trends for ten administrative regions, excluding Addis region. Survey indicator fixed effects allow separate intercepts for districts sampled only in NLFS, or only in DHS. Time-invariant controls: terrain ruggedness, log distance to zone capital, log distance to wholesale argricultural market, log distance to Addis, soil quality, district area. Initial values of time-varying controls: log population density, nighttime luminosity, temperature, standard deviation of precipitation, built-up area and literacy. Windmeijer conditional F denotes the Sanderson and Windmeijer (2016) weak instrument F-statistic for multiple endogenous variables, reported separately for the two first stages (columns (1) and (2)).

To check the robustness of the core reduced-form results discussed above, I briefly highlight further results regarding education and the labour force composition. With respect to education, the overall education results (Table S11) are ambiguous since only road-connected districts show increases in literacy, whereas educational attainment is insignificant. Labour force participation (Table S12) increases insignificantly from roads, and decreases strongly with additional electrification, in line with the above narrowing of the age pyramid: more teenagers are out of the labour force.

Table A2: Second Stage: Occup. [ISIC] on Roads & Elec., NLFS sample (1999-2013)

	Agr. [isic]	Ser. [isic]	Man. [isic]
	(1)	(2)	(3)
Road Indicator	−0.020	0.072	−0.050
	(0.089)	(0.068)	(0.041)
Road*Elec Ind.	−0.216	0.084	0.138*
	(0.150)	(0.111)	(0.071)
Sample	*NLFS*	*NLFS*	*NLFS*
Year FE	✓	✓	✓
Region-Year Trend	✓	✓	✓
Controls	✓	✓	✓
Windmeijer cond. F.	27.095	9.156	
p-val $\beta_1 + \beta_2 = 0$	0.090	0.115	0.221
p-val $\beta_1 - \beta_2 = 0$	0.338	0.937	0.040
R^2	0.277	0.321	−0.032
Adjusted R^2	0.252	0.297	−0.068
Observations	1163	1163	1163

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. All standard errors clustered at district-level. Sample includes all districts sampled at least twice in NLFS (1999, 2005, 2013), yielding a total of 1163 district-years. Districts in Addis Abeba and Somali regions are excluded, the latter due to partial omission and non-random sampling of Somali pastoralists in NLFS. Individual answers to survey question about 'current occupation' grouped into sectors according to ISIC industry classifications. 'Agriculture' denotes aggregated share of district-year survey sample of working age that respond to work in an agricultural occupation. Omitting results for outcome variable 'Currently Not Working', which denotes share of survey sample of working age not currently working. Road indicator and road indicator–electricity indicator interaction denote predicted values from the two first stage specifications. Region-year-specific time trends for nine administrative regions, excluding Addis and Somali regions. Time-invariant controls: terrain ruggedness, log distance to zone capital, log distance to wholesale argricultural market, log distance to Addis, soil quality, district area. Initial values of time-varying controls: log population density, nighttime luminosity, temperature, standard deviation of precipitation, built-up area and literacy. Windmeijer conditional F denotes the Sanderson and Windmeijer (2016) weak instrument F-statistic for multiple endogenous variables, reported separately for the two first stages (columns (1) and (2)).

Table A3: Second Stage: Occup. on Roads & Elec., NLFS balanced panel (1999-2013)

	Agr.		Ser.		Man.	
	(1)	(2)	(3)	(4)	(5)	(6)
Road Indicator	-0.046	-0.021	0.134^{**}	0.103^{**}	-0.083^{*}	-0.081^{*}
	(0.086)	(0.072)	(0.063)	(0.053)	(0.047)	(0.042)
Road*Elec Ind.	-0.194	-0.254^{*}	0.106	0.167	0.094	0.084
	(0.144)	(0.135)	(0.104)	(0.103)	(0.074)	(0.066)
Sample	*NLFS*	*NLFS (bal.)*	*NLFS*	*NLFS (bal.)*	*NLFS*	*NLFS (bal.)*
Year FE	✓	✓	✓	✓	✓	✓
Region-Year Trend	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
Windmeijer cond. F.	27.095	36.91	9.156	10.978		
p-val $\beta_1 + \beta_2 = 0$	0.061	0.042	0.008	0.007	0.879	0.967
p-val $\beta_1 - \beta_2 = 0$	0.456	0.168	0.845	0.617	0.084	0.056
R^2	0.282	0.301	0.231	0.254	0.132	0.193
Adjusted R^2	0.257	0.271	0.205	0.223	0.102	0.159
Observations	1163	866	1163	866	1163	866

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. All standard errors clustered at district-level. 'Balanced' sample includes all districts sampled at least twice across all three NLFS rounds (1999, 2005, 2013), but only using the earliest and latest district-year observation available, yielding a total of 866 district-years. Districts in Addis Abeba and Somali regions are excluded, the latter due to partial omission and non-random sampling of Somali pastoralists in NLFS.

Table A4: Gender: Occup. on Roads & Elec., NLFS sample (1999-2013)

	Female			Male		
	Agr.	Ser.	Man.	Agr.	Ser.	Man.
Road Indicator	-0.036	0.204^{**}	-0.166^{*}	-0.077	0.086	-0.000
	(0.113)	(0.082)	(0.085)	(0.078)	(0.056)	(0.032)
Road*Elec Ind.	-0.207	0.150	0.076	-0.146	0.059	0.084
	(0.181)	(0.138)	(0.124)	(0.137)	(0.095)	(0.059)
Sample	*NLFS*	*NLFS*	*NLFS*	*NLFS*	*NLFS*	*NLFS*
Year FE	✓	✓	✓	✓	✓	✓
Region-Year Trend	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓
Windmeijer cond. F.	27.601	9.219		26.505	8.915	
p-val $\beta_1 + \beta_2 = 0$	0.132	0.003	0.424	0.063	0.074	0.136
p-val $\beta_1 - \beta_2 = 0$	0.504	0.779	0.180	0.712	0.841	0.274
R^2	0.253	0.096	0.190	0.313	0.296	0.268
Adjusted R^2	0.227	0.064	0.162	0.289	0.271	0.242
Observations	1163	1163	1163	1163	1163	1163

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. All standard errors clustered at district-level. NLFS female and male respondents aggregated separately to district-year samples aggregated separately for women (columns 1-3) and men (columns 4-6). Specification and sample otherwise as in Supplement Table S7.

Table A5: Second Stage: Migration on Roads & Elec., NLFS sample (1999-2013)

	Mig.<1yr	Mig.<2yr	Mig.<6yr	Mig. ever
	(1)	(2)	(3)	(4)
Road Indicator	0.001	−0.007	−0.009	−0.022
	(0.011)	(0.017)	(0.038)	(0.080)
Road*Elec Ind.	0.027	0.043	0.119*	0.244*
	(0.020)	(0.029)	(0.066)	(0.143)
Sample	*NLFS*	*NLFS*	*NLFS*	*NLFS*
Year FE	✓	✓	✓	✓
Region-Year Trend	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Windmeijer cond. F.	27.095	9.156		
p-val $\beta_1 + \beta_2 = 0$	0.105	0.144	0.062	0.084
p-val $\beta_1 - \beta_2 = 0$	0.325	0.211	0.151	0.169
R^2	0.067	0.079	0.146	0.267
Adjusted R^2	0.034	0.047	0.116	0.241
Observations	1163	1163	1163	1163

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. All standard errors clustered at district-level. Sample includes all districts sampled at least twice in NLFS (1999, 2005, 2013), yielding a total of 1163 district-years. Districts in Addis Abeba and Somali regions are excluded, the latter due to partial omission and non-random sampling of Somali pastoralists in NLFS. Individual answers to survey question about 'did you migrate in last one/two/six years' (columns (1), (2), (3)) or 'did you ever migrate?' (column (4)). Road indicator and road indicator–electricity indicator interaction denote predicted values from the two first stage specifications. Region-year-specific time trends for nine administrative regions, excluding Addis and Somali regions. Time-invariant controls: terrain ruggedness, log distance to zone capital, log distance to wholesale argricultural market, log distance to Addis, soil quality, district area. Initial values of time-varying controls: log population density, nighttime luminosity, temperature, standard deviation of precipitation, built-up area and literacy. Windmeijer conditional F denotes the Sanderson and Windmeijer (2016) weak instrument F-statistic for multiple endogenous variables, reported separately for the two first stages (columns (1) and (2)).

Table A6: Real Consumption: Wealth, Assets and Housing Changes on Roads and Elec., DHS sample (2000-2016)

	Wealth Index	Radio	TV	Refrig.	Bike	Scooter	Car	Phone	Elec.	Tap Water	Toilet	Floor
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Road Indicator	0.086	0.011	−0.097**	−0.065	−0.029	0.001	−0.004	−0.006	−0.046	−0.331**	−0.029	−0.049
	(0.362)	(0.061)	(0.042)	(0.060)	(0.018)	(0.003)	(0.004)	(0.016)	(0.075)	(0.161)	(0.023)	(0.050)
Road*Elec Ind.	0.928*	0.052	0.139	0.066	0.030	−0.002	0.011	0.029	0.302*	0.645**	0.056	0.068
	(0.556)	(0.105)	(0.086)	(0.048)	(0.023)	(0.011)	(0.008)	(0.031)	(0.172)	(0.267)	(0.058)	(0.103)
Sample	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*	*DHS*
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Region-Year Trend	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Windmeijer cond. F.	19.344	7.885										
p-val $\beta_1 + \beta_2 = 0$	0.058	0.508	0.558	0.977	0.968	0.944	0.223	0.443	0.119	0.198	0.541	0.832
p-val $\beta_1 - \beta_2 = 0$	0.275	0.770	0.038	0.172	0.100	0.791	0.160	0.360	0.094	0.008	0.259	0.384
R^2	0.555	0.112	0.287	0.117	0.018	0.008	0.011	0.093	0.379	0.091	0.034	0.357
Adjusted R^2	0.555	0.111	0.286	0.116	0.017	0.007	0.010	0.092	0.379	0.090	0.033	0.356
Observations	46425	46414	46411	35617	46409	46407	46409	46406	46410	46421	46409	46409

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$. All standard errors clustered at district-level. Sample includes individual household heads' answers from (enumeration areas in) districts sampled at least twice in DHS (2000, 2005, 2011, 2016), yielding a total of 46425 individual-year observations. Respondents in Addis Abeba and Somali regions are excluded. Household heads' answers to survey questions about asset ownership and housing characteristics. Wealth index in quantiles provided by DHS. Road indicator and road indicator–electricity indicator interaction denote predicted values from the two first stage specifications. Region-year-specific time trends for ten administrative regions, excluding Addis region. Time-invariant controls: terrain ruggedness, log distance to zone capital, log distance to wholesale agricultural market, log distance to Addis, soil quality, district area. Initial values of time-varying controls: log population density, nighttime luminosity, temperature, standard deviation of precipitation, built-up area and literacy. Windmeijer conditional F denotes the Sanderson and Windmeijer (2016) weak instrument F-statistic for multiple endogenous variables, reported separately for the two first stages (columns (1) and (2)).

A.VII Appendix: Alternative Model Setup

The main version of the model presented in Section 5 features constant returns to scale production technology in all sectors, with infrastructure complementarities arising solely from non-homotheticities on the demand side. An alternative model setup that gives rise to qualitatively similar infrastructure complementarities can be arrived at by assuming increasing returns to scale in the production of electricity-using sectors. In the following, such an alternative model setup is presented and briefly discussed.

The three main differences to the main model in Section 5 are: first, sectoral demand is assumed to be homothetic; second, the number of manufacturing varieties produced in a given location is assumed to be endogenous and location-specific; third, production technology in electricity-using sectors is assumed to exhibit increasing returns to scale of varying strength depending on each location's electrification status.

Regarding the first, sectoral composite goods are assumed to exhibit standard homothetic demand (Michaels et al., 2011) such that the original equation (6) with $\gamma^K = 0$ across all sectors $K \in \{T, M, S\}$ becomes:

$$C_n = \left[\sum^{K} \psi^K \left(C_n^K \right)^{\frac{\kappa-1}{\kappa}} \right]^{\frac{\kappa}{\kappa-1}} \tag{20}$$

while maintaining that sectoral composite goods are complementary, $0 < \kappa < 1$.

Regarding the second, manufacturing varieties are assumed to be endogenous, with each location producing its own measure of varieties, M_i. Demand for manufacturing varieties in equation (7) becomes a standard Helpman (1998) expression:

$$C_n^M = \left[\sum_{i \in N} \int_0^{M_i} \left(c_{ni}^M(j) \right)^{\nu} dj \right]^{\frac{1}{\nu}} \tag{21}$$

where I maintain the assumption that varieties within the manufacturing sector are substitutes, $\sigma = \frac{1}{1-\nu} > 1$. To avoid confusion with the total number of locations N in equation (21), M denotes the manufacturing sector, whereas M_i denotes location i's endogenous number of varieties j. The associated Dixit-Stiglitz price index for the manufacturing sector (cf. Supplement S.II) is subject to an equivalent change and prices will also be integrated over an endogenous number of varieties.

Regarding the third, a standard assumption to obtain increasing returns to scale in production is Krugman's (1991) fixed cost in production: to produce a manufacturing variety j, a firm in location i must incur fixed cost F_i (in units of labour l_i), in addition to the variable costs.[100] The labour demand per variety j in location i is thus: $l_i(j) = F_i + \frac{y_i(j)}{A_i}$. Zero-profit and downward-sloping demand imply that output across varieties is constant in a given location and only depends on productivity, fixed cost and the elasticity of substitution, $y_i(j) = \bar{y}_i = A_i(\sigma - 1)F_i$.[101] In equilibrium, the endogenous number of varieties, M_i is proportional to labour supply L_i: $M_i = \frac{L_i}{\sigma F_i}$.

Recent evidence on last-mile electrification provides a microfoundation for the crucial question how increasing returns to scale may vary with electrification: Figueiredo Walter and Moneke (2023) document new evidence on sizeable local fixed costs to grid connection that firms in low-income countries face upon arrival of the electric grid. These fixed costs arise from the technical requirement to erect a local trunk network for distribution. Therefore, the share of end-users that adopt electricity conditional on grid arrival is bimodal in many low-income countries, including Ethiopia (Figueiredo Walter & Moneke, 2023) – with a large mass at zero adoption, and sizeable mass at near-full adoption – implying real indivisibilities for firms to get grid-connected.

Such large fixed cost to make productive use of electricity can be modelled by assuming that they differ across locations as a function of local grid arrival, $elec_i$:[102]

$$F_i(elec_i) = \begin{cases} F & \text{if } elec_i = 0 \\ F + E & \text{if } elec_i = 1 \end{cases}$$

The above formulation implies stronger increasing returns in electrified locations, which implies starker agglomeration patterns from electrification than without it. An advantage of this modelling choice is that the elasticity of substitution would govern the strength of the additional economies of scale from electrification. A mild compli-

[100] An alternative assumption would be one of external economies of scale, i.e. making manufacturing sector productivity in electrified locations a function of labour employed in this sector. Although such an alternative assumption would have different implications for prices in general equilibrium, it would have qualitatively similar implications for the presence of an infrastructure complementarity.

[101] For clarity of exposition, I present a simple case where firms only require one input (labour) instead of two (land and labour). The model's qualitative predictions are unchanged with two inputs.

[102] For simplicity, I assume that the fixed cost of connecting to the grid E do not vary across locations, whereas overall fixed costs F_i will, due to locations' varying electrification status $elec_i$.

cation, in turn, is that an appropriate value for F has to be estimated from Ethiopian microdata, for example using the CSA's Large- and Medium-Scale Manufacturing Surveys. The key parameter of interest, E, can be estimated in analogue to the main model's estimation of ϕ, except that electricity can now affect both fixed (E) and variable cost (ϕ).[103,104] I note that this model is also amenable to quantitative analysis in a similar fashion to the analysis in Section 6, which I leave for future work.

A.VIII Appendix: Additional Structural Estimation Results

In the main model, the construction of new all-weather roads is assumed to affect trade cost by decreasing the effective least-cost distance between potential trading locations across Ethiopia. Table A7 provides a sanity check of this assumption, confirming that the employed procedure to arrive at least-cost distances (explained in greater detail in Supplement S.III) is meaningful. Reassuringly, when benchmarked against state-of-the-art routing engines, distances appear sensible: random district-pair distances from this matrix are very close to distances predicted by, e.g., OpenRouteService.

Table A7: Full-panel: Roads and Least-Cost Distances (2000-2016)

	Dependent variable:			
	Log(Sum of Least-cost Distances)			
	(1)	(2)	(3)	(4)
Roads Ind.	-0.203^{***}	-0.114^{***}	-0.045^{***}	-0.040^{***}
	(0.011)	(0.007)	(0.009)	(0.003)
Controls		✓	✓	
Year FE			✓	✓
District FE				✓
Observations	2,752	2,744	2,744	2,752
Note:		*p<0.1; **p<0.05; ***p<0.01		

[103] Electrification improves a location's variable cost of manufacturing (via increasing productivity at rate ϕ), while worsening fixed costs F (via adding E). Thus, to match the reduced-form result of growing manufacturing employment in electrified locations, an upper bound on E is required.

[104] A related complication arises from the fact that the structural estimation procedure would need to estimate two parameters (E and ϕ) from one equation (over N locations). Hence, to separately identify both parameters, additional restrictions need to be placed on the estimation procedure.