Where Urban Agglomerations Grow Next

Build-Vintage as a Map for Service and Investment

Evidence from 26 Global-South Countries

Karan Singh Bagavathinathan*

Independent Economist

July 2026

Abstract

A city's economic boundary, its contiguous built-up agglomeration, routinely outgrows the boundary that governs it. Dating every pixel of urban fabric in 27 Global-South countries by the decade it first urbanized (GHSL build-vintage), we split each agglomeration into a pre-1990 core and three decadal expansion rings and place georeferenced DHS households on the rings in 26 of them across two survey waves (334,091 households, 758 agglomerations). We report three findings. First, within the same city, deprivation falls with the age of the fabric; households on the 2010s ring are 0.37 SD poorer than their own core, with 10 to 12 pp less flush sanitation and clean cooking, while the 1990s ring is at the core. Second, following fixed rings across two waves, these deficits narrow on average, so the frontier is usually a transient cohort the city grows out of, but absorption takes about a decade and the delay is costly, in schooling gaps that persist, out-of-pocket spending on missing services, hours lost commuting, and private investment kept off unserviced land, so it drags on the whole city's productivity and not only the households on it. Third, catch-up is not universal; Dhaka and Dushanbe stay deprived across both waves and Amman inverts the pattern through refugee densification of its old core, while which cities diverge tracks neither region nor income and development-bank finance is neither necessary nor sufficient for closing the gap. Because DHS frames predate the newest fabric they under-cover the frontier, so the estimates lean on the better-sampled cities and if anything understate the deficit. How fast a city closes its frontier gap, and why a few never do, is the causal question this design sets up.

Keywords: urban agglomeration, build-vintage, remote sensing, urban frontier, service delivery, convergence, development finance, Global South.
JEL classification: O18, R11, R14, H54, I32, O15.

*Preliminary working paper. Comments welcome at vbk.singh@gmail.com.

1 Introduction

In economic theory a country's cities are its most productive asset. Across much of the Global South, though, their expansion is pushed by rural poverty, climate stress, and conflict as much as it is pulled by opportunity. A city runs on labour across the income range, so the open question is where its lower-income residents settle and whether services reach them there. Vienna, with its large stock of serviced social housing, shows they can be housed well and centrally, and it is one of several cities that run public housing at this scale. In much of the Global South they settle the edge instead, where fabric is built *before* it is governed and before water, sewerage, and tenure arrive.

The frontier stays unserviced for two reasons. The first is cost. Connecting a new ring to trunk water, power, and roads costs more the farther it sits from the network. Then there is the energy to pump and run it. Few public budgets can fund that ahead of settlement, or subsidize the first arrivals, so the frontier waits. The fiscal squeeze behind this is long-standing and structural, rooted in the rising real cost of the labour-intensive services a city runs against stretched budgets. Some cities cap outward growth on purpose. Portland's urban growth boundary and the green belts around London and Seoul push building upward instead of out (Brueckner, 2000; Bertaud, 2018), and many emerging economies write the same rules into law, from China's urban growth boundaries and farmland lines to India's master-plan green belts and Cape Town's urban edge. A boundary contains sprawl only if the city can house the demand it pushes inward, at higher density and serviced. That needs enforcement to hold the line and public money to fund the housing and the trunk capacity behind it. Where either is missing, informal building is the cheaper option, so the boundary is crossed and the fabric fills in unserviced.

The second reason is governance. A single agglomeration often spans many municipalities, so the trunk lines that would serve it cross several budgets and no one body owns the whole. Coordinating a project across them is hard, and the smaller municipalities cannot raise their share or buy into a larger scheme (Ahrend et al., 2017; Bahl et al., 2013). A few countries have built metropolitan authorities to get around this, by amalgamating councils in South Africa or layering a metro government over London, but these reforms take a higher tier to impose and stay rare; the common outcome is a weak coordinating body that can plan more than it can compel (Bird and Slack, 2007).

External finance has tried to fill the gap. Development banks and donors fund frontier water, sanitation, and upgrading, from output-based subsidies for household connections to sites-and-services and low-cost condominial sewers (Michaels et al., 2021). The recurring problem is scale. Reaching the mass of the frontier rather than a pilot district is where these efforts fall short. The Bank once tried to reach it household by household, through sites-and-services schemes meant to recover their cost from the settlers who took the plots; collection seldom held, the public budget absorbed the shortfall, and the approach was wound down as lending shifted to city governments (Buckley and Kalarickal, 2006). That frontier has only grown since, and it still needs to be reached. What failed was the way the money was recovered, not the goal, so what would help now is a way to aim it, to show which rings to service first. Section 5 returns to this with our multi-lender evidence.

Whether the frontier is absorbed into serviced fabric or hardens into a permanent under-serviced periphery is central to urban policy, and each of the two data sources sees only half of it. Satellite data see the fabric but not the people; survey data see the people but not the age of the fabric under them. Individual studies have joined them within a single city. What has been missing is one measure applied the same way across many countries, so the size of the frontier's deprivation can be compared rather than described case by case.

This paper links the two. We organize households by when the land under them was built. Using the GHSL built-up surface at four epochs (1990/2000/2010/2020), we define urban agglomerations (UAs) morphologically, as contiguous built-up clusters $\geq$10 km^2 independent of administrative boundaries, and date each pixel by its first built epoch, yielding a pre-1990 *core* and decadal *expansion rings*. Georeferenced DHS clusters are assigned to rings; DHS households then reveal each ring's character. Informality is therefore an *empirical outcome*. A ring counts as informal exactly when its households are deprived.

We report three findings. First, within the same city, deprivation falls with the age of the fabric, so households on the newest ring are the poorest and least serviced, across 26 countries. Second, following fixed rings across two survey waves, the gap closes in most cities, but slowly, over about a decade, and that lag is itself the cost that matters. The frontier cohort's schooling deficit outlasts the servicing gap, households spend out of pocket to substitute for the missing network and time to commute around it, and land that would draw private investment stays out of productive use, so the unserviced edge holds down the city's productivity and its delivery of public goods to all residents as much as it burdens the households on it. In a minority of cities, Dhaka and Dushanbe among them, the gap does not close at all. Third, development-bank finance reaches some frontiers and not others and is neither necessary nor sufficient for catch-up; banks also tend to lend where local capacity is already strong, so the link is a correlation we cannot yet read as cause.

2 Related literature

Measuring the city as its competitive edge. Governments, and the development banks that still lend to the poorer among them, increasingly judge a city by its economic performance, its productivity, its connectivity, and the services that underpin them, and invest on that basis (World Bank, 2009, 2015). Who pays varies with income. The lowest-income countries lean on external finance, while middle-income states run large urban programs from their own budgets; India, now a middle-income borrower rather than an aid recipient, funds national missions such as Smart Cities and AMRUT itself. Either way, measuring the city has become part of measuring its competitive edge, and two data families do that measuring and rarely meet. Satellite and remote-sensing work captures the fabric, how far a city has spread, in what form, and when it grew (Angel et al., 2016; Pesaresi et al., 2023; Dijkstra et al., 2021), with economic delineations of the functional city drawn from lights, buildings, or commuting (Dingel et al., 2021; de Bellefon et al., 2021; Baragwanath et al., 2021) and city shape itself shown to matter for welfare (Harari, 2020). Household surveys

of the DHS type capture the people, the gradients in wealth, water, sanitation, and health across urban space (Fink et al., 2014; Gollin et al., 2021; Lagakos, 2020; Henderson et al., 2021), and the economics of slums (Marx et al., 2013; Fox, 2014). Each sees half. The satellite sees the ground but not the household standing on it; the survey sees the household but not the decade its ground was built. A competitiveness assessment can therefore score a city's spread and its average deprivation and still not say whether the growing edge is being absorbed or left behind. To our knowledge, and consistent with a systematic search, no prior study links the build-date of urban fabric to household welfare across countries. Dating the fabric under the household is the measure that gap needs, and it is this paper's central addition.

The urban agglomeration outgrows the governed one. A city's built-up area expands past the boundary that administers it. The monocentric model explains the outward push (Brueckner, 2000), yet it says little about who governs the new edge. In much of the Global South one agglomeration spans many municipalities, and the smaller peripheral ones can neither raise their share of a trunk project nor buy into a larger scheme (Ahrend et al., 2017; Bahl et al., 2013). Metropolitan authorities that could coordinate across them are rare and costly to impose, so the common arrangement leaves the periphery under a body with thin revenue and little hold over the trunk networks that cross it (Bird and Slack, 2007). This literature is fiscal and institutional, measured at the level of the city or the council, and it rarely reaches the household on the actual expansion ring. A vintage measure locates where that governance gap lands in household welfare, on the fabric that grew up outside the body meant to service it.

What has been tried to reach the frontier. There is a long doctrine for servicing expansion before it settles. The planned-extension approach holds that reserving an arterial grid and trunk rights-of-way ahead of growth keeps later servicing cheap (Angel et al., 2016; Bertaud, 2018). Development banks once funded this directly through sites-and-services, and the evidence is that planned frontier infrastructure produces durably better neighborhoods (Michaels et al., 2021). The approach was wound down when its costs did not come back from the households it served (Buckley and Kalarickal, 2006), and attention moved to upgrading fabric already built, where gains are partial and sometimes reversed (Harari and Wong, 2025), and to titling, which matters though it does not build the network (Field, 2007; Galiani and Schargrodsky, 2010). The rich-country benchmark for all of this is filtering, where new fabric houses the better-off and old fabric filters down (Rosenthal, 2014; Brueckner and Rosenthal, 2009). These studies work city by city. What has been missing is a cross-country panel of where unplanned frontiers close their deficit and where they harden, and a common way to place each city against that filtering benchmark. We build both.

Why the frontier is expensive, and whether that changes. The cost of servicing the edge is not a fixed number. Part of it is structural. The labour-intensive services a city runs resist productivity gains, so their real cost rises against stretched budgets, the cost disease Baumol tied to the urban crisis (Baumol, 1967; Baumol et al., 1985). Part of it is geometric. Fragmented, leapfrog

growth and rugged peri-urban terrain make trunk water, sewerage, and drainage run farther and climb more to reach the same households (Harari, 2020; Bertaud, 2018; Nunn and Puga, 2012), and the flood-prone low ground that the frontier poor often settle raises both the cost of protecting a network and the odds it is damaged (Kocornik-Mina et al., 2020; Rentschler et al., 2022). Whether technology lowers the bill is the open question. Households already substitute where a market exists, buying vended or packaged water when the pipe does not arrive, and off-grid power and on-site sanitation relax some margins. The binding deficits in our data are the networked services, piped supply and sewerage, which have no cheap decentralized substitute yet, so technology thins the problem at the edges and leaves the capital-heavy trunk in place. Work on developing-city transit (Tsivanidis, 2026; Balboni et al., 2020; González-Navarro and Turner, 2018; Jedwab and Moradi, 2016) and on geocoded aid (Dreher et al., 2022; Isaksson and Kotsadam, 2018) supplies the tools to study where finance lands, though it does not interact the money with the expansion stage of the fabric it serves. That interaction, finance against build-vintage, is what Section 5 supplies.

3 Data and method

Fabric. We use the GHSL GHS-BUILT-S R2023A built-up surface at 100 m resolution for the epochs 1990, 2000, 2010, and 2020. A pixel counts as built when at least 20% of its area is built-up, and an agglomeration is an 8-connected cluster of at least 10 km^2, clipped to national boundaries (geoBoundaries). Each pixel's vintage is the first epoch in which it was built. The pixel-level definition handles polycentric and merged agglomerations automatically; mergers are 5 to 11% of agglomerations but 24 to 66% of built-up footprint.

Households. We link the DHS geographic, household, and individual recodes (GE, HR, and IR/MR) across 26 countries at two waves each, 52 surveys in all. Each survey cluster is assigned to the build-vintage ring that dominates within a 2 km buffer, a radius that also absorbs the DHS geocode displacement (up to 2 km for urban clusters and 5 km for rural ones; Burgert et al., 2013). All estimates use DHS sampling weights. Occupation and employment variables were read from each survey's own value labels rather than a common scheme, and five surveys that shipped broken or non-standard codes were corrected.

Estimation. Regressions are run within agglomerations, with standard errors clustered at the DHS cluster level and household wealth standardized within each country-survey. Per-country tables and the coding audit are provided in an online appendix.

The sample narrows in two steps. We build build-vintage rasters for 27 countries (Table 2). Twenty-six enter the household panel (Table 1). Madagascar drops out, not for want of surveys (it has two georeferenced waves) but because its DHS clusters fall almost entirely in the pre-1990 core, leaving no expansion-ring households to measure, even though over a quarter of its fabric is post-1990. Four of the twenty-six (Armenia, Benin, Cameroon, Haiti) were added most recently, and their agglomerations are compact, so their DHS clusters fall overwhelmingly in the pre-1990 core: together they contribute only about thirty newest-ring clusters. Whether that reflects

genuinely little recent expansion or a survey frame that misses the newest edge, it is the same under-coverage the paper is careful about throughout, and it leaves the frontier estimates resting on the better-sampled countries. Within the panel we test catch-up as a within-UA ring×late interaction (Table 6), which two waves estimate only weakly.

It is worth being explicit about coverage. The sample is agglomerations, so we keep contiguous built-up clusters of at least 10 km^2, and smaller towns and isolated built-up patches fall below the threshold and are left out. The paper therefore speaks to a country's larger cities rather than to every urban settlement. The rule is uniform across countries, and the number of agglomerations it returns tracks how urbanized each country is, from 2 in Lesotho and Armenia to 316 in India. Across the 26 panel countries it defines 770 agglomerations, of which 758 contain georeferenced DHS households and enter the estimation. The per-country counts are the UAs column of Table 1. Of the 758, 708 are observed in both waves and 50 in only one, and the median UA contributes about 129 households per wave (range 18 to 4,430).

Table 1: The household panel: 26 countries, two DHS waves each. "Largest UA" names each country's biggest agglomeration for orientation, while the panel uses every UA in the country (India's 316 span Delhi, Mumbai, and the other metros). "Survey years" gives the two DHS waves. UAs = agglomerations with in-sample households.

Country	Largest UA	Survey years	UAs	Clusters	Households
India	Delhi	2015–16, 2019–21	316	7,733	157,920
Nigeria	Lagos	2018, 2024	118	848	24,625
Jordan	Amman	2017, 2023	5	606	12,259
Angola	Luanda	2015, 2023	21	435	11,149
Philippines	Manila	2017, 2022	18	473	9,635
Bangladesh	Dhaka	2017, 2022	17	256	9,357
Benin	Cotonou	2012, 2017	11	357	8,478
Mozambique	Maputo	2011, 2022	28	379	8,293
Ghana	Accra	2014, 2022	23	286	8,033
Cameroon	Douala	2011, 2018	13	299	7,401
DR Congo	Kinshasa	2013, 2023	32	218	7,367
Kenya	Nairobi	2014, 2022	11	314	6,812
Zambia	Lusaka	2018, 2024	20	270	6,480
Tajikistan	Dushanbe	2017, 2023	13	296	6,383
Burkina Faso	Ouagadougou	2010, 2021	10	234	5,916
Haiti	Port-au-Prince	2012, 2016	3	197	5,779
Côte d'Ivoire	Abidjan	2012, 2021	16	219	5,680
Tanzania	Dar es Salaam	2015, 2022	25	240	5,368
Ethiopia	Addis Ababa	2019, 2024	13	189	5,099
Mali	Bamako	2018, 2023	13	182	4,487
Armenia	Yerevan	2010, 2015–16	2	149	3,558
Senegal	Dakar	2019, 2023	11	168	3,471
Malawi	Lilongwe	2015, 2024	6	102	2,967
Cambodia	Phnom Penh	2014, 2021	5	111	2,846
Lesotho	Maseru	2014, 2023	2	106	2,538
Nepal	Kathmandu	2016, 2022	6	77	2,190
26 countries		**52 surveys**	**758**	**14,744**	**334,091**

4 Results

4.1 Expansion is large and recent

Table 2: Built-up footprint (km^2) by first-built epoch, ranked by total footprint. Core = fabric present by 1990; 90s/00s/10s = area first built in that decade.

#	Country	Total	Core	90s	00s	10s	#	Country	Total	Core	90s	00s	10s
1	India	11,607	5,869	2,640	2,099	1,000	15	Benin	447	245	45	109	48
2	Nigeria	6,099	3,080	526	1,647	845	16	Mali	434	238	77	69	49
3	Ghana	1,516	724	281	354	157	17	Senegal	433	255	48	71	59
4	DR Congo	1,329	884	112	134	200	18	Kenya	397	254	59	48	36
5	Philippines	1,130	845	190	65	31	19	Tajikistan	351	177	90	63	21
6	Angola	1,008	707	85	68	149	20	Jordan	337	142	86	68	41
7	Tanzania	999	571	153	156	120	21	Haiti	198	107	30	24	37
8	Mozambique	901	442	140	199	122	22	Nepal	196	95	41	40	20
9	Zambia	627	412	88	89	37	23	Cambodia	185	28	49	78	30
10	Bangladesh	615	372	121	96	26	24	Malawi	173	128	14	18	13
11	Cameroon	535	347	61	80	47	25	Madagascar	171	122	28	15	6
12	Côte d'Ivoire	530	340	49	68	73	26	Armenia	135	103	18	10	5
13	Burkina Faso	483	224	85	120	54	27	Lesotho	91	45	23	17	6
14	Ethiopia	477	231	79	107	60							

Expansion in these countries is large and recent. About half of the 2020 built-up footprint was added after 1990 (a median of 45% across countries; Figure 5, right panel), and in the fastest cases the post-1990 build exceeds the pre-1990 core. Phnom Penh is 85% new fabric, and Amman, Ouagadougou, Accra, Kathmandu, and Addis Ababa are all above half. The magnitude is concentrated, with India and Nigeria alone holding 56% of the sample's urban fabric (left panel).

4.2 The vintage gradient (within-city)

Table 3 reports these within-city regressions, one column per outcome, each coefficient the ring's gap to its own city's pre-1990 core. The gradient holds across all four outcomes and steepens on newer fabric. On the 2010s ring the gaps in flush sanitation, electricity, and clean cooking are all close, roughly 10 to 12 percentage points, and the wealth gradient, at about 0.37 SD, is steeper still, while the 1990s ring is indistinguishable from the core in wealth and sanitation. No single service escapes the deficit.

The gradient is in people as much as pipes. Human capital shows a similar frontier deficit. In a within-city regression of women's schooling on ring dummies (women 15–49, latest wave, 26 countries, single-year age controls, UA fixed effects; Table 4), the 2000s and 2010s rings have significantly less schooling than their own city's core, about 0.7 to 0.8 fewer years and 4 to 5 percentage points lower secondary completion, while the 1990s ring is close to the core. The shape differs from wealth in one way: the schooling deficit is largest on the 2000s ring rather than the newest, and it survives age controls, so it reflects a genuine cohort settled on the frontier and not

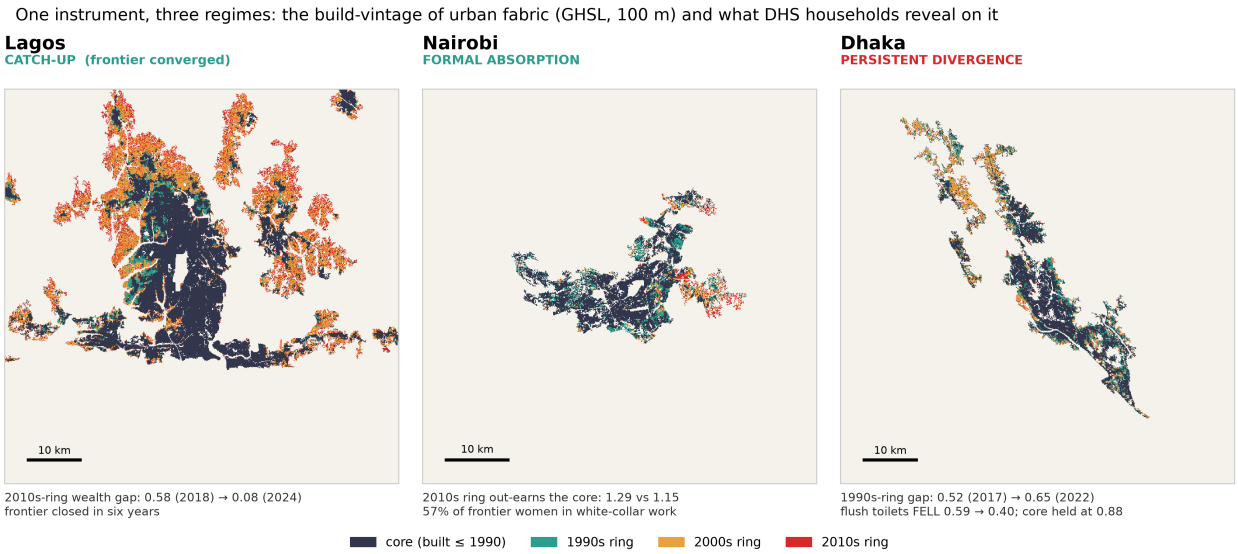

Figure 1: One instrument, three regimes. Build-vintage of urban fabric (GHSL, 100 m) for Lagos, Nairobi, and Dhaka: dark = core built by 1990; teal/amber/red = 1990s/2000s/2010s expansion rings. Lagos has a vast recent fringe, Nairobi a compact one, and Dhaka's river-constrained fabric leaves little room to expand, so its rings *densify* instead and services are outrun. The morphology shows why Dhaka is the hard case.

Table 3: Within-city vintage gradient. Each column is a separate regression of the outcome on the three ring dummies with UA fixed effects (latest wave, one cross-section, so no time dummy; pre-1990 core is the omitted category).

	(1) Wealth (z)	(2) Flush toilet	(3) Electricity	(4) Clean cooking
1990s ring	−0.032	−0.006	−0.011**	−0.022**
	(0.041)	(0.011)	(0.005)	(0.009)
2000s ring	−0.210***	−0.050***	−0.037***	−0.039***
	(0.032)	(0.011)	(0.008)	(0.008)
2010s ring	−0.372***	−0.112***	−0.113***	−0.111***
	(0.087)	(0.018)	(0.024)	(0.031)
UA fixed effects	Yes	Yes	Yes	Yes
Agglomerations	745	745	745	745
Observations	168,132	168,132	168,132	168,132

Notes: coefficients are the gap to the core; standard errors clustered by DHS cluster in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

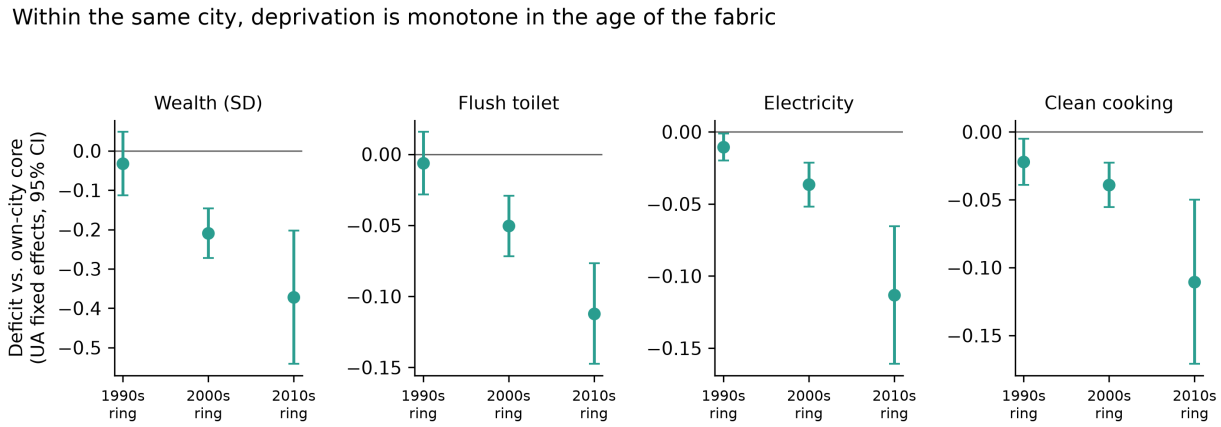

Figure 2: The vintage gradient: within-agglomeration coefficients (Table 3) with 95% CIs. Deprivation rises on newer fabric; the 1990s ring is statistically at the core in wealth and sanitation, with only small gaps in electricity and clean cooking.

just younger arrivals. The frontier deficit is therefore more than infrastructure that can be laid in; it is also a cohort with less schooling, which closes more slowly.

Table 4: Schooling gradient, within-city. Women 15–49, latest wave; each column is a separate within-UA regression of the outcome on ring dummies with single-year age dummies and UA fixed effects (pre-1990 core omitted).

	(1) Years of schooling	(2) Secondary completion
1990s ring	−0.241	−0.015*
	(0.150)	(0.009)
2000s ring	−0.817***	−0.050***
	(0.110)	(0.008)
2010s ring	−0.706***	−0.046***
	(0.202)	(0.014)
Age dummies	Yes	Yes
UA fixed effects	Yes	Yes
Agglomerations	745	745
Observations	183,248	183,251

Notes: coefficients are the gap to the core; standard errors clustered by DHS cluster in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.3 The frontier ring vs. the core, country by country

Table 5 takes this country by country, comparing the newest ring's DHS-weighted mean with its own core in the latest wave. In most countries the newest ring is the poorer and less-serviced side of the pair, often by a wide margin. Ghana's 2010s ring sits about 1.6 SD below its core in wealth and has no flush sanitation against the core's 58 percent, and Mozambique, Tanzania, Angola, and

Table 5: The newest expansion ring vs. the core, country by country, latest DHS wave. Levels are DHS-weighted means pooled over each country's agglomerations.

		Wealth		Flush			
Country (survey)	Newest ring	Core	Ring	Core	Ring	Frontier	Clusters
Ghana (2022)	2010–20	1.01	−0.63	0.58	0.00	poorer	3
Senegal (2023)	2010–20	1.06	0.50	0.89	0.60	poorer	2
Mozambique (2022)	2010–20	1.10	0.53	0.18	0.03	poorer	11
Angola (2023)	2010–20	1.09	0.58	0.69	0.44	poorer	15
Tanzania (2022)	2010–20	1.09	0.72	0.51	0.37	poorer	8
India (2019–21)	2010–20	9.89	4.84	0.91	0.77	poorer	51
Cambodia (2021)	2010–20	1.53	0.81	1.00	0.93	poorer	2
Burkina Faso (2021)	2010–20	1.27	0.91	0.11	0.13	poorer	4
Bangladesh (2022)	2000–10	0.91	0.35	0.88	0.88	poorer	8
Tajikistan (2023)	2000–10	0.72	−0.46	0.60	0.09	poorer	13
Zambia (2024)	2000–10	1.08	0.88	0.52	0.40	poorer	15
Benin (2017)	2000–10	1.01	0.81	0.19	0.15	poorer	23
Nigeria (2024)	2010–20	0.83	0.75	0.79	0.78	poorer	14
DR Congo (2023)	2010–20	1.93	1.71	0.42	0.18	poorer	3
Malawi (2024)	2010–20	1.40	0.93	0.19	0.10	poorer	2
Côte d'Ivoire (2021)	2010–20	1.11	0.84	0.73	0.62	poorer	3
Haiti (2016)	2010–20	1.21	0.65	0.36	0.18	poorer	3
Ethiopia (2024)	2010–20	1.40	1.26	0.33	0.32	thin	5
Kenya (2022)	2010–20	1.15	1.29	0.76	0.73	thin	10
Nepal (2022)	2010–20	1.65	1.48	0.99	0.98	thin	2
Mali (2023)	2010–20	1.13	1.21	0.24	0.19	thin	2
Cameroon (2018)	2000–10	0.86	0.95	0.35	0.38	thin	3
Armenia (2015–16)	1990–2000	0.53	0.61	0.95	1.00	thin	3
Jordan (2023)	2010–20	0.42	0.87	0.97	0.91	richer	10
Philippines (2022)	1990–2000	0.69	0.57	1.00	0.98	similar	30

Notes: levels are DHS-weighted means pooled over each country's agglomerations, so this is a within-country core-vs-ring comparison; the within-city gradient with UA fixed effects is Table 3. "Newest ring" is the most recent ring with adequate clusters, which is 2000–10 for Bangladesh, Tajikistan, Zambia, Benin, and Cameroon and 1990–2000 for the Philippines and Armenia, where later fabric is too thin to read. "Clusters" counts the newest ring's DHS clusters; rows with two or three are indicative only. "Frontier" marks whether the newest ring is poorer than, similar to, or richer than the core in wealth; "thin" marks rings with too few clusters per city, or a within-city split, to sign reliably. A within-UA check confirms only Jordan's inversion is robust and the Philippines is genuinely at parity; the frontier signs for Kenya, Mali, Ethiopia, Nepal, Cameroon, and Armenia do not hold up city by city, and Nigeria's gap is small because Lagos has all but closed. Wealth is z-scored within country-survey except India (raw NFHS index), so compare within a row, not across. Jordan's core is *poorer* than its ring (refugee-densified old Amman), the MENA inversion; Cambodia's better-sampled 2000s ring is converging (Table 6); Lesotho and Madagascar have no ring clusters in-sample.

Bangladesh run the same way. Two cases break the pattern, and both survive a within-city check. Amman's refugee-densified core is poorer than its ring, and Manila's older frontier has reached parity. The rings marked thin rest on two or three clusters and are indicative only, and because wealth is standardized within each country-survey (India's raw index excepted), the levels compare within a row, not across.

4.4 The two-wave dynamics are underpowered

Table 6: Catch-up: within-city regression across both waves. Ring dummies are the earlier-wave gap to the core; the Ring×Late interactions are the change in the later wave, so a positive interaction means the ring's deficit narrows. UA and wave fixed effects; SEs clustered by DHS cluster.

	(1) Wealth (z)	(2) Flush	(3) Electricity	(4) Clean cooking
1990s ring	-0.087^{*}	-0.028^{**}	-0.004	-0.030^{***}
	(0.045)	(0.011)	(0.006)	(0.011)
2000s ring	-0.253^{***}	-0.068^{***}	-0.035^{***}	-0.071^{***}
	(0.035)	(0.011)	(0.007)	(0.010)
2010s ring	-0.591^{***}	-0.163^{***}	-0.141^{***}	-0.132^{***}
	(0.087)	(0.025)	(0.032)	(0.020)
1990s × Late	0.038	0.000	-0.010^{*}	-0.015
	(0.054)	(0.014)	(0.006)	(0.013)
2000s × Late	0.030	0.021	-0.006	0.039^{***}
	(0.042)	(0.014)	(0.009)	(0.012)
2010s × Late	0.233^{**}	0.060^{**}	0.053	0.041
	(0.116)	(0.030)	(0.041)	(0.034)
UA fixed effects	Yes	Yes	Yes	Yes
Wave fixed effects	Yes	Yes	Yes	Yes
Observations	334,091	334,071	321,828	328,011

Notes: the ring dummies reproduce the vintage gradient. The Ring×Late interactions are the direct two-wave catch-up, and they are small and mixed: positive and significant at 5% for the newest ring in wealth and sanitation, positive for the 2000s ring in clean cooking, and near zero or slightly negative for electricity. Two adjacent surveys give little power to detect a ring closing; the stronger evidence for catch-up is the cross-section (Table 3). SEs clustered by DHS cluster. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The direct two-wave catch-up is weak. The Ring×Late interactions in Table 6 are small and mixed across indicators: positive and significant at 5% for the newest ring in wealth and sanitation, positive and significant for the 2000s ring in clean cooking, and near zero for electricity. Two adjacent surveys, four to eleven years apart, give little power to watch a ring close. The stronger evidence for catch-up is the cross-section itself. In Table 3 the 1990s ring's gap to the core is already gone, its coefficient near zero, while the 2010s ring sits well below the core, so a ring narrows its deficit as its fabric ages, over a span longer than our two waves can capture.

A constant observation window sharpens the frontier estimate. The two waves are four to eleven years apart, and pooling them lets a binary later-wave dummy stand for very different

amounts of elapsed time. Restricting to countries whose waves are the same distance apart removes that mixing without assuming the gap closes at a constant annual rate. On a common five-to-seven-year window (14 countries) and a tighter five-to-six-year window (11 countries), the newest ring's narrowing is larger than in the full pool and significant for sanitation, electricity, and clean cooking, with wealth marginally significant on the tighter window (Table 7). The sub-sample still contains Dhaka and Dushanbe, so the net catch-up holds with the two non-closing cities left in. The estimate is larger but rests on fewer countries, so we read it as corroboration of the frontier's tendency to close, not as a sharper measurement of the rate. Splitting instead by window length points the same way on the wealth index, where the gap closes more in the countries observed longer (about 0.46 SD over nine years against 0.19 over five, at a similar annual pace), but the two bands are different countries and the service-by-service pattern moves around. What the two waves support, then, is the average tendency of the frontier to close; the differences on top of it, how fast any one city closes and how large the two divergences are, are real in the data but beyond what two waves and thin frontier samples can pin down.

Table 7: Newest-ring catch-up on a constant observation window. Each cell is the 2010s ring × Late interaction from the Table 6 specification (within-UA and wave fixed effects, SEs clustered by DHS cluster), re-estimated on countries whose two DHS waves are the same distance apart, so a common window replaces the mix of 4-to-11-year gaps.

2010s ring × Late	Full sample (26 countries)	5–7 yr window (14 countries)	5–6 yr window (11 countries)
Wealth (z)	0.233**	0.334	0.458*
	(0.116)	(0.262)	(0.270)
Flush toilet	0.060**	0.093	0.105*
	(0.030)	(0.058)	(0.062)
Electricity	0.053	0.211**	0.177
	(0.041)	(0.102)	(0.110)
Clean cooking	0.041	0.097**	0.108**
	(0.034)	(0.045)	(0.047)

Notes: positive = the newest ring's deficit narrows between waves. The five-to-six year sample is Armenia, Bangladesh, Benin, Ethiopia, Jordan, Mali, Nepal, Nigeria, Philippines, Tajikistan, and Zambia, and it includes the two non-closing cities. Point estimates exceed the full sample because a common window no longer averages over unequal 4–11 year gaps; the cost is fewer countries and wider confidence intervals. $^{*}p < 0.10$, $^{**}p < 0.05$.

Catch-up in the pipes does not reset human capital. A ring can regain water and power within a decade, but the schooling gap does not close with the infrastructure. The frontier is also a cohort with 0.7 to 0.8 fewer years of schooling (Table 4), and that residue outlives the servicing deficit that produced it.

The same two-wave limit applies to the two cities whose frontiers do not close. Dhaka's and Dushanbe's rings stay wide or widen while other cities' narrow. Both are capital-primate systems where ring densification outruns networked services, under universal electricity (Soviet-era, in Dushanbe's case). With only two such cities observed over two waves, we cannot say how large the gap is, only that these two are the exception to the frontier's tendency to close.

4.5 Frontier character is a city-level regime

Table 8: Expansion-regime typology. The serviced-frontier cases are the ones that survive a within-city check; most others are deprived to varying degrees, a few are too thinly sampled to sign, and two cities do not close.

Type	Countries	Signature
Serviced frontier	Jordan (Amman, inverted), Philippines (Manila, absorbed)	at or above the core; robust within-city
Deprived frontier (the common case)	Nigeria, Ghana, India, Tanzania, Mozambique, Angola, Zambia, Burkina Faso, Cambodia, DR Congo, Senegal, Benin, Côte d'Ivoire, Haiti; Ethiopia, Nepal, Malawi (thin)	newest ring poorer than the core; older rings less so
Frontier not signed (thin)	Kenya, Mali, Cameroon, Armenia	at or above the core but too thin to sign
Persistent, non-closing	Bangladesh (Dhaka), Tajikistan (Dushanbe)	deprived, does not close over two waves
No measurable frontier	Madagascar, Lesotho	expanded, but no ring clusters sampled

Frontier character varies city by city, not just by region or income. Amman inverts the pattern entirely, a poorer core than frontier from refugee densification of old Amman, and it is the one inversion robust within-city; Manila's older frontier has been fully absorbed. Against those, Dhaka and Dushanbe are the two cities where the frontier does not close. The labour data track welfare. Deprived frontiers are agrarian-to-urban transition zones, where frontier men's agriculture share runs four to six times the core's in India, Ghana, and Tanzania. India's frontier is unique in one dimension. Women's labour-force participation there falls to about 9%, against 24 to 28% in every other Indian zone, and no such fall appears elsewhere, including Dhaka's garment belt, where ring women out-participate the core.

5 Development-bank finance and the frontier

We scan the World Bank Projects API for urban, water, sanitation, and transport commitments $\geq$\$10M since 1995 in the study countries and keep those whose project titles are genuinely urban (511 projects). Several cities show sustained agglomeration-targeted lending (Ouagadougou's ONEA water-utility sequence 2009–2024; Accra's Greater Accra sanitation and water programme, GAMA; Kenya's Urban Support Programs; Dar es Salaam's metropolitan development project, DMDP, and its bus rapid transit; Luanda's \$500M Bita water project, which serves the *southern expansion*; Kinshasa's \$500M Kin-Elenda urban project). India's metros lead the sample at roughly \$8B of World Bank urban lending. The two non-closing cities anchor the range. Dushanbe is barely financed, with about \$0.14B, the smallest urban portfolio in the sample. Dhaka has one of the larger portfolios, about \$2.5B from the World Bank (and more from the ADB and JICA), yet its large operations target the core network, other cities, or rural districts while its rings densify.

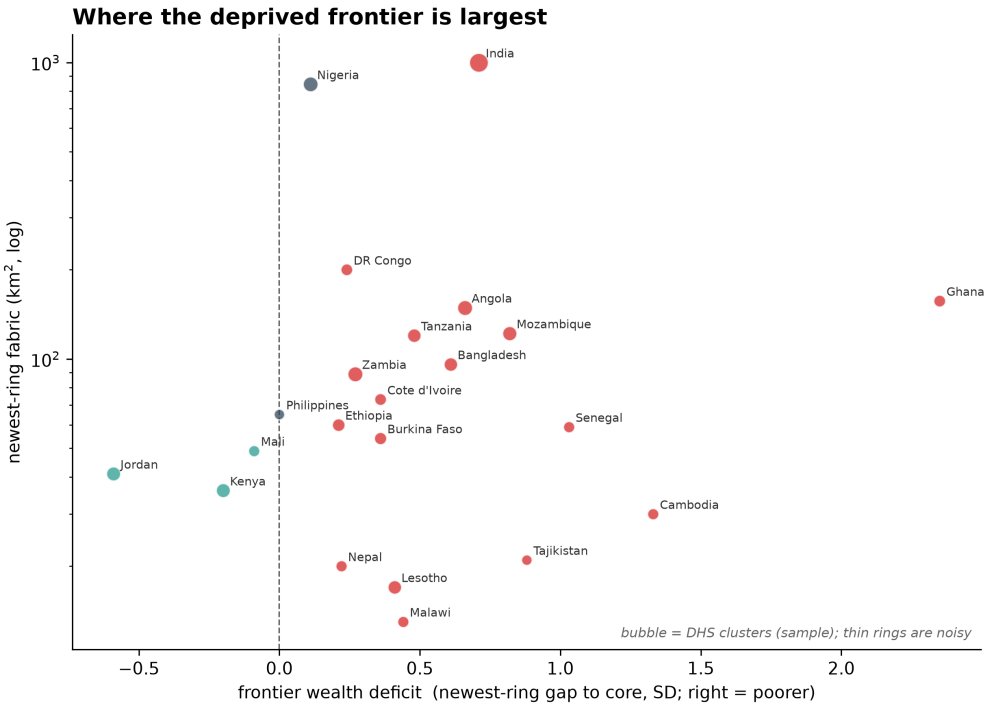

Figure 3: Where the deprived frontier is largest. Each point is one country. For that country's newest ring, the horizontal axis is its wealth deficit to the core (right = poorer; negative = a richer frontier, as in Amman), the vertical axis is the ring's fabric area, and bubble size is DHS clusters, so thin rings appear small. The large, deprived, well-sampled frontiers, India, Nigeria, Mozambique, and Tanzania, are the targeting priority; Ghana's far-right point rests on three clusters.

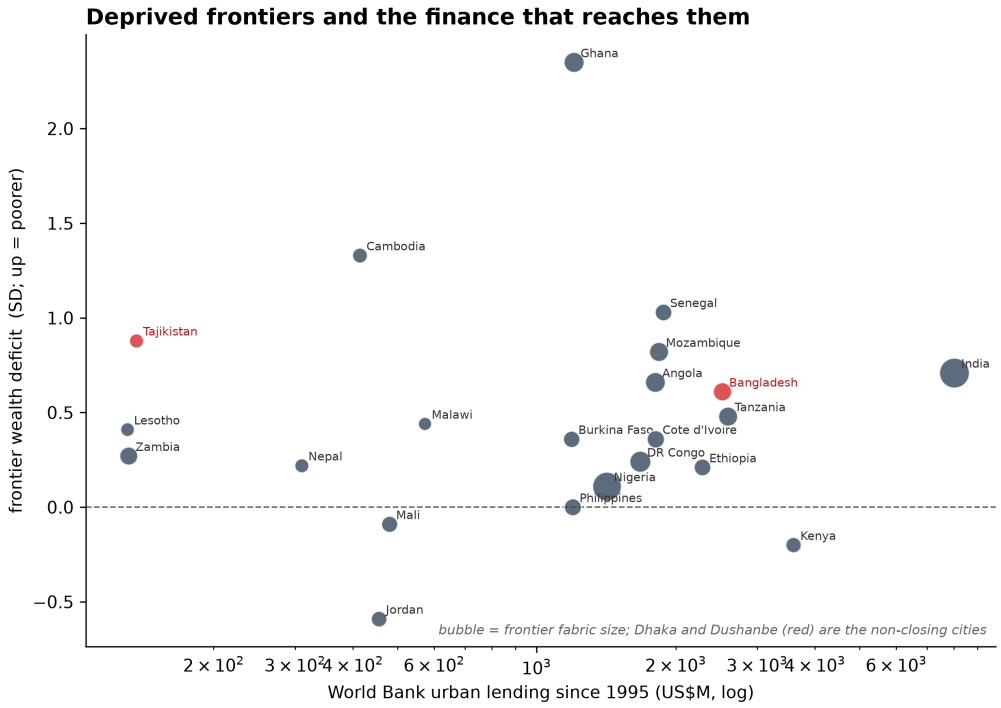

Figure 4: Deprived frontiers and the World Bank urban finance that reaches them. Each point is one country. The horizontal axis is World Bank urban lending since 1995 (log) and the vertical axis is the frontier's wealth deficit. Tajikistan (red) carries a large deficit on the smallest urban portfolio (about $0.14B), the barely-financed case; Bangladesh (red) carries a large deficit despite one of the larger portfolios (about $2.5B), the financed-but-not-reaching case. Bubble size is frontier fabric size. We plot World Bank lending only; the non-Bank figures are reliable as individual projects, not as portfolio totals (see text).

Beyond the World Bank. The Bank is not the only lender, and in several cities it is not the largest. To compare lenders on one basis we use OECD Creditor Reporting System commitments for water, sanitation, and urban development since 1995, which cover the World Bank, the AfDB, and France/AFD on common purpose codes; the EIB does not report to this system and is added from its own project data. Two settings stand out. In Amman the World Bank is only about three percent of urban lending. The EIB leads, with over €1.5B for Jordan's water sector since the 1980s (the framework loans behind Amman's supply and the Aqaba–Amman desalination and conveyance line), and France/AFD adds about $0.8B. Amman is a special case, since Jordan is chronically water-scarce and hosts large refugee populations, so much of this water finance answers that pressure rather than frontier expansion, and its inverted gradient is refugee densification of the old core, not a frontier that closed. The more general pattern is francophone West Africa, where the AfDB and AFD together exceed the World Bank in Dakar ($0.79B against $0.60B), Abidjan ($0.76B against $0.71B), and Ouagadougou ($0.60B against $0.57B). The clearest single case is Dakar's KMS3 water project (€433M), only about one percent World Bank, with AFD (€108M), the Islamic Development Bank (€133M), the EIB (€100M), and the AfDB (€65M) carrying the rest. Urban transport reinforces the pattern, since Abidjan's metro and Dakar's bus rapid transit are largely French and multi-lender financed. In the anglophone and Asian cities, by contrast, the World Bank still leads.

Whether any lender deliberately targets the frontier ring, rather than the core or the city at large, is a question the project data cannot answer, because a project names a city but not the ring inside it, and making that link needs the geocoded overlay reserved for later work. The per-city World Bank magnitudes above come from our project scan (water, sanitation, urban development, and urban transport, by project title); the cross-lender comparison uses OECD CRS purpose codes, which isolate water, sanitation, and urban development but not urban transport, with the EIB added from its own reporting since it is absent from the CRS. All figures are gross commitments since 1995 in current dollars.

Does finance cause catch-up, or does the Bank choose the best areas? The correlation cannot distinguish the two, and selection runs in both directions. IDA country allocation follows a formula weighted by the CPIA governance score, which selects on the same institutional capacity that drives ring servicing, and within countries banks lend where a competent implementer already exists (Ouagadougou's ONEA). The non-Bank lenders discussed above cover the AfDB, EIB, AFD, and EBRD, but Chinese, ADB, and JICA finance remain out, so "unfinanced" labels stay weakest for the Asian cities (Dhaka's ADB network projects and JICA metro in particular). Two facts in our own data nonetheless bound the interpretation. Dhaka shows that finance is not sufficient, since one of the larger urban portfolios sits on a diverging system, and Phnom Penh and Lilongwe show that it is not necessary, since both converge with little World Bank lending. What matters appears to be whether money reaches the ring and meets a governance structure able to absorb it. Identifying the causal contribution requires three things: (i) geocoded project sites overlaid on the vintage rasters,

giving finance exposure by ring, with UA×wave fixed effects so that only ring-siting selection survives; (ii) pre-trends from third DHS waves (available for Ghana, Kenya, Senegal, Tanzania, and others), testing whether treated rings were already converging before flagship projects (GAMA 2013, DMDP 2015, Bita 2019); and (iii) instruments for presence, including IDA allocation-formula and graduation shocks and the post-2015 SDG-6 sector-wide WASH surge. This is the design of the causal follow-on.

6 Mechanisms: why the frontier deficit is an economic cost

The deficit is more than an equity problem for the households on the frontier; it is a drag on the productivity of the city as a whole, and that reframes it for a growth audience. Three channels connect the two, and our household data carries the footprint of each. *First, unconnected density is not productive density.* The gains that make cities productive, namely matching, sharing, and learning (Duranton and Puga, 2004), require density that is actually connected to trunk power, water, and roads. A frontier that adds population but not connective infrastructure raises headcount without raising the effective economic density that agglomeration economies run on. *Second, spatial mismatch taxes both sides of the labour market.* Workers settle the cheap frontier while jobs remain in the core; absent transit, the commute eats the wage (Bertaud, 2018; Tsivanidis, 2026; Harari, 2020). Firms face a thinner effective labour pool and workers a lower effective wage at the same time. *Third, an informality equilibrium with an exploitation edge.* Where power, water, and secure tenure never arrive, formal firms do not locate, so residents work informally, either self-employed or in the low-wage market within reach. Our labour tables show this household footprint. The newer the ring, the more manual and informal the work mix and the less white-collar. Weak outside options combined with commuting costs hand employers monopsony-like power, which is the concrete form "exploitation" takes, namely low pay, no contract, and no protection, because the worker cannot credibly walk (Marx et al., 2013; Henderson et al., 2021). On top of this sits a regressive private tax. Where the network fails, households buy around it, which is our packaged-water result, and this deepens the squeeze that makes low-wage work harder to refuse. Over time the schooling gap documented above caps the frontier cohort's productivity, so today's neglect compounds into tomorrow's lower city output. The implication for a development lender is that the frontier deficit belongs under growth and competitiveness as much as under poverty. We stress that our evidence is household-side; the firm-level link is drawn from the agglomeration and spatial-mismatch literature, so this section is mechanism and motivation rather than a firm-level test.

7 Policy implications

(1) The frontier deficit is, in most cities, a *timing lag* with a catch-up window of roughly a decade. Policy's job is to shorten it, and build-date gives a satellite-observable, near-real-time instrument for where the deficit is *now*. The instrument does not depend on the source of funds. A finance

ministry funding its own urban mission, a city budgeting its trunk extensions, and a development bank choosing where to lend can all read it the same way. (2) The deficit spans every service we measure, and the margins households cannot buy their way around are piped water and sanitation, so those are where public provision matters most. (3) The frontier is an agrarian-to-urban transition zone, so urban governance and the tax base should be extended before informal consolidation sets in. (4) The two cities whose frontiers do not close, Dhaka and Dushanbe, are the exception to the catch-up tendency rather than a large estimated penalty, and their contrasting financing histories point to different remedies, presence where urban money is largely absent (Dushanbe) and targeting where it is present but core-directed (Dhaka). (5) Whether infrastructure that *leads* expansion prevents entry into the failure regime is the open causal question, left to a governance event study and a finance-and-transit design. (6) Part of the frontier's growth is distress-driven rather than pulled by opportunity. In Sub-Saharan Africa rainfall failure and drought push rural households off failing land and toward cities (Barrios et al., 2006; Henderson et al., 2017), and violent conflict forces rural-to-urban displacement (Ibáñez and Vélez, 2008; Depetris-Chauvin and Santos, 2018). Policy that helps rural households survive these shocks in place, drought-resilient agriculture, crop insurance, and safety nets, together with resolution of the conflicts that drive displacement, would ease the distress share of frontier inflow, upstream of and complementary to servicing it. We offer this as a hypothesis rather than a lever this paper tests, and flag that the wider link from rural investment to urban growth is ambiguous, since raising farm productivity can release labour to cities as much as retain it (Gollin et al., 2014; Emerick, 2018; Bryan et al., 2014). The defensible case is for shock-survival support, not for rural investment in general.

8 Limitations and next steps

Our identification here is descriptive. Sorting into rings is differenced out only insofar as it is time-invariant, so the causal layer is left to future work.

The sample is conservative about the frontier itself. DHS sampling frames are drawn from censuses that predate the newest fabric, so the most recent and least formal settlements are the ones most likely to be under-covered, and some are missed altogether. If anything this understates the frontier deficit, because the worst-served ground is the hardest to enumerate. Several countries also have thin newest-ring samples, flagged per row; the best-powered anchors are Nigeria, Burkina Faso, Cambodia, Mozambique, Tajikistan, India, Kenya, and Benin. India alone is 316 agglomerations and about two-fifths of the households, so we confirm that both the within-city gradient and the catch-up hold with India dropped (Appendix A).

The finance results are correlational, and the correlation is hard to read, because the cities that attract programmes and lending tend also to have the capacity to service a frontier, so money and governance arrive together. Separating them is the next step. We plan to take specific projects with known start dates, the Accra, Dar es Salaam, and Luanda water and sanitation programmes, together with a third DHS survey wave, and to check whether the rings those projects reached

were already improving before the money arrived. If they were, the finance was following progress rather than causing it. A parallel design uses census-town reclassification to isolate the effect of being governed as urban.

A label-driven corrected water variable (packaged water classified from each survey's own codes) shows packaged water is an urban-*core* substitution phenomenon (74% of in-UA households in Ghana and the Philippines; 52% Amman; 46% Kathmandu; zero in Dushanbe, whose rings cannot buy their way out). Once packaged sources are counted, "any improved water" is near-universal and flat across rings, which confirms that the binding frontier margins are piped *network* service and, above all, sanitation, and that frontier households bear a hidden retail cost of water.

GHSL 2010/2020 epochs partly interpolate satellite anchors (2014/2018); an anchor-year robustness re-cut is planned.

A Robustness to dropping India

India contributes 316 of the sample's agglomerations and about two-fifths of its households, so we re-estimate the two within-city specifications with India excluded (Table 9). The vintage gradient, the 2010s ring's gap to the core, stays negative and significant on all four outcomes, and the two-wave catch-up, the 2010s ring × Late interaction, is if anything larger. Neither headline result rests on India. Dropping it does shrink the wealth gradient in levels, because India's raw wealth index has a wide spread that widens the level gap, but the sign and significance hold throughout, and the service gradients and the catch-up estimates move very little.

References

Ahrend, R., E. Farchy, I. Kaplanis, A.C. Lembcke. What makes cities more productive? Evidence from five OECD countries on the role of urban governance. *J. Regional Science* 57(3), 385–410.

Bahl, R., J. Linn, D. Wetzel (eds.). *Financing Metropolitan Governments in Developing Countries*. Lincoln Institute of Land Policy.

Bird, R., E. Slack. An approach to metropolitan governance and finance. *Environment and Planning C: Government and Policy*.

Buckley, R., J. Kalarickal (eds.). *Thirty Years of World Bank Shelter Lending: What Have We Learned?* Directions in Development, World Bank.

Nunn, N., D. Puga. Ruggedness: the blessing of bad geography in Africa. *Review of Economics and Statistics*.

Kocornik-Mina, A., T. K. J. McDermott, G. Michaels, F. Rauch. Flooded cities. *American Economic Journal: Applied Economics*.

Rentschler, J., M. Salhab, B. A. Jafino. Flood exposure and poverty in 188 countries. *Nature Communications*.

World Bank. *World Development Report 2009: Reshaping Economic Geography*. World Bank.

World Bank. *Competitive Cities for Jobs and Growth: What, Who, and How*. World Bank.

Table 9: Robustness to dropping India. Newest-ring (2010s) coefficients from the gradient specification (Table 3) and the catch-up specification (Table 6), full sample against India excluded. Within-UA fixed effects (and wave effects in the catch-up rows), DHS weights, SEs clustered by DHS cluster.

	Full (26 countries)	India excluded (25)
Panel A. Vintage gradient: 2010s ring gap to core		
Wealth (z)	-0.372^{***}	-0.211^{**}
	(0.087)	(0.092)
Flush toilet	-0.112^{***}	-0.104^{***}
	(0.018)	(0.021)
Electricity	-0.113^{***}	-0.181^{***}
	(0.024)	(0.035)
Clean cooking	-0.111^{***}	-0.053^{***}
	(0.031)	(0.018)
Panel B. Catch-up: 2010s ring × Late		
Wealth (z)	0.233^{**}	0.399^{**}
	(0.116)	(0.156)
Flush toilet	0.060^{**}	0.048
	(0.030)	(0.036)
Electricity	0.053	0.148^{**}
	(0.041)	(0.064)
Clean cooking	0.041	0.056^{*}
	(0.034)	(0.029)
Agglomerations (A / B)	745 / 758	433 / 442

Notes: Panel A is the latest-wave cross-section; Panel B pools both waves. A negative Panel A coefficient is a frontier deficit; a positive Panel B coefficient is the deficit narrowing between waves. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

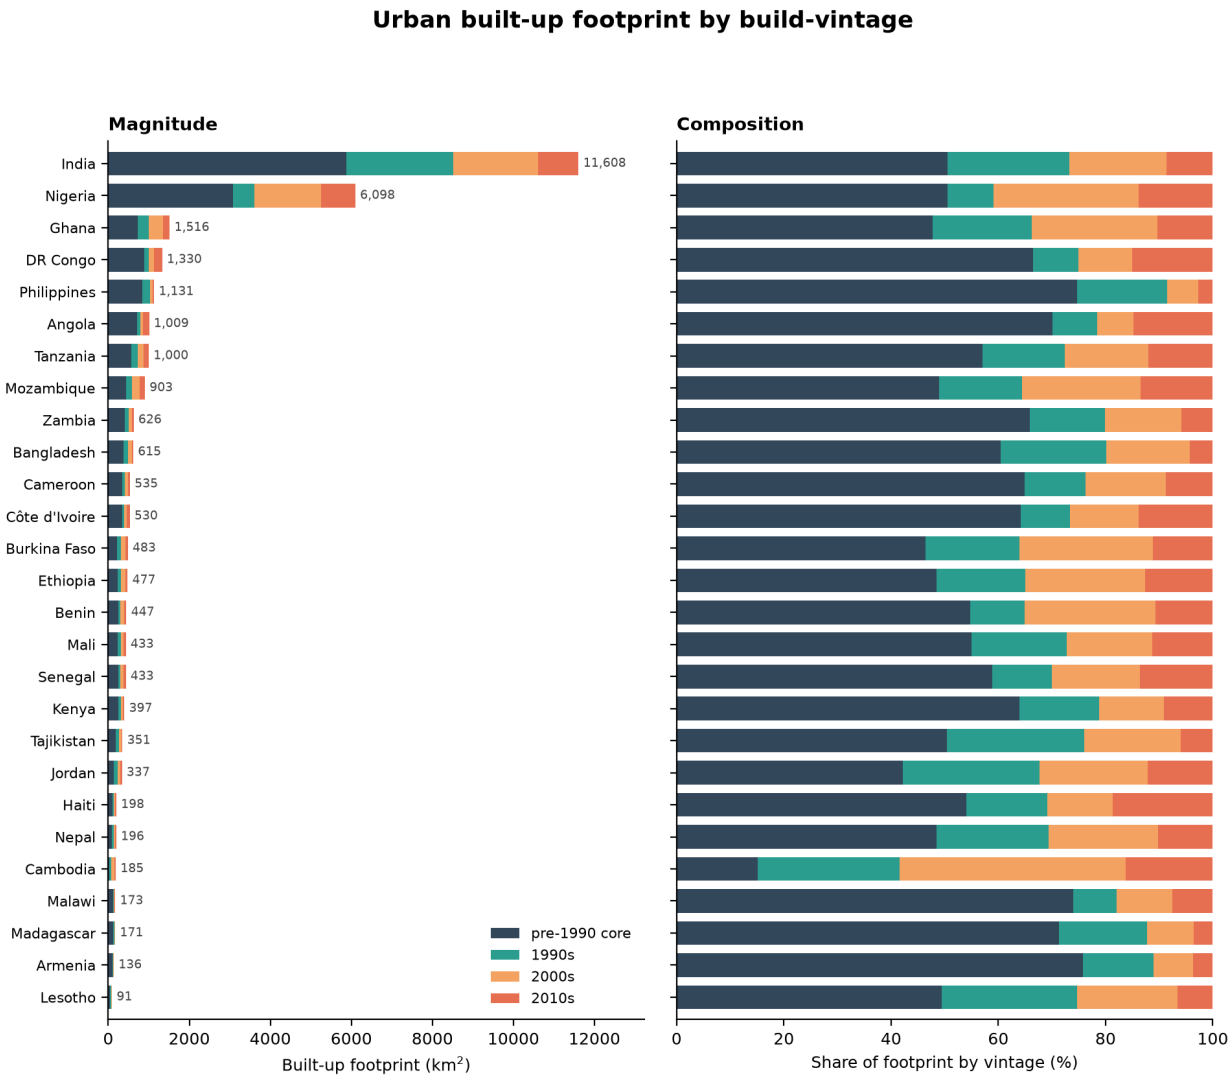

Figure 5: Built-up footprint by build-vintage, the same data as Table 2. Left panel: magnitude in km^2, stacked by the decade each area was first built. Right panel: the same countries as 100% shares, so the post-1990 expansion is legible regardless of city size. Countries are ranked by total footprint.

Brueckner, J. Urban sprawl: diagnosis and remedies. *International Regional Science Review*.

Angel, S., J. Parent, D. Civco, A. Blei. *Atlas of Urban Expansion*. Lincoln Institute / UN-Habitat / NYU.

Balboni, C., G. Bryan, M. Morten, B. Siddiqi. Transportation, gentrification, and urban mobility: the inequality effects of place-based policies. Working paper.

Barrios, S., L. Bertinelli, E. Strobl. Climatic change and rural–urban migration: the case of sub-Saharan Africa. *J. Urban Economics* 60(3), 357–371.

Bryan, G., S. Chowdhury, A.M. Mobarak. Under-investment in a profitable technology: the case of seasonal migration in Bangladesh. *Econometrica* 82(5), 1671–1748.

Baumol, W.J. Macroeconomics of unbalanced growth: the anatomy of urban crisis. *American Economic Review* 57(3).

Baumol, W.J., S.A.B. Blackman, E.N. Wolff. Unbalanced growth revisited: asymptotic stagnancy and new evidence. *American Economic Review* 75(4).

Bertaud, A. *Order without Design: How Markets Shape Cities*. MIT Press.

Duranton, G., D. Puga. Micro-foundations of urban agglomeration economies. In *Handbook of Regional and Urban Economics*, vol. 4.

Baragwanath, K., R. Goldblatt, G. Hanson, A. Khandelwal. Detecting urban markets with satellite imagery: an application to India. *J. Urban Economics*.

de Bellefon, M.-P., P.-P. Combes, G. Duranton, L. Gobillon, C. Gorin. Delineating urban areas using building density. *J. Urban Economics*.

Brueckner, J., S. Rosenthal. Gentrification and neighborhood housing cycles. *Review of Economics and Statistics*.

Burgert, C., J. Colston, T. Roy, B. Zachary. Geographic displacement procedure and georeferenced data release policy for the DHS. *DHS Spatial Analysis Reports 7*.

Dijkstra, L., A. Florczyk, S. Freire, T. Kemper, M. Melchiorri, M. Pesaresi, M. Schiavina. Applying the degree of urbanisation to the globe. *J. Urban Economics*.

Dingel, J., A. Miscio, D. Davis. Cities, lights, and skills in developing economies. *J. Urban Economics*.

Depetris-Chauvin, E., R.J. Santos. Unexpected guests: the impact of internal displacement inflows on rental prices in Colombian host cities. *J. Development Economics* 134, 289–309.

Dreher, A., A. Fuchs, B. Parks, A. Strange, M. Tierney. *Banking on Beijing*. Cambridge University Press.

Emerick, K. Agricultural productivity and the sectoral reallocation of labor in rural India. *J. Development Economics* 135, 488–503.

Field, E. Entitled to work: urban property rights and labor supply in Peru. *Quarterly J. Economics*.

Fink, G., I. Günther, K. Hill. Slum residence and child health in developing countries. *Demography*.

Fox, S. The political economy of slums: theory and evidence from Sub-Saharan Africa. *World Development*.

Galiani, S., E. Schargrodsky. Property rights for the poor: effects of land titling. *J. Public Economics*.

Gollin, D., D. Lagakos, M.E. Waugh. The agricultural productivity gap. *Quarterly J. Economics* 129(2), 939–993.

Gollin, D., M. Kirchberger, D. Lagakos. Do urban wage premia reflect lower amenities? Evidence from Africa. *J. Urban Economics* 121, 103301.

González-Navarro, M., M. Turner. Subways and urban growth: evidence from Earth. *J. Urban Economics*.

Harari, M. Cities in bad shape: urban geometry in India. *American Economic Review*.

Harari, M., M. Wong. Slum upgrading and long-run urban development: evidence from Indonesia. *Review of Economic Studies*, advance access.

Henderson, J.V., T. Regan, A. Venables. Building the city: from slums to a modern metropolis. *Review of Economic Studies*.

Henderson, J.V., A. Storeygard, U. Deichmann. Has climate change driven urbanization in Africa? *J. Development Economics* 124, 60–82.

Ibáñez, A.M., C.E. Vélez. Civil conflict and forced migration: the micro determinants and welfare losses of displacement in Colombia. *World Development* 36(4), 659–676.

Isaksson, A.-S., A. Kotsadam. Chinese aid and local corruption. *J. Public Economics*.

Jedwab, R., A. Moradi. The permanent effects of transportation revolutions in poor countries: evidence from Africa. *Review of Economics and Statistics*.

Lagakos, D. Urban–rural gaps in the developing world. *J. Economic Perspectives*.

Marx, B., T. Stoker, T. Suri. The economics of slums in the developing world. *J. Economic Perspectives*.

Michaels, G., D. Nigmatulina, F. Rauch, T. Regan, N. Baruah, A. Dahlstrand. Planning ahead for better neighborhoods: long-run evidence from Tanzania. *J. Political Economy*.

Pesaresi, M., et al. GHS-BUILT-S R2023A: built-up surface grid 1975–2030. *JRC, European Commission*.

Rosenthal, S. Are private markets and filtering a viable source of low-income housing? *American Economic Review*.

Tsivanidis, N. Evaluating the impact of urban transit infrastructure: evidence from Bogotá's TransMilenio. *American Economic Review* 116(2), 418–463.