Rerouting Behavior and Travel-Time Penalties Under Flood-Induced Road Closures

Ridwan Tiamiyu[1,*]; Oladimeji Basit Alaka[2]; Oluwabusola Ajayi[3]
[1]Fountain Construction Company, Nigeria
[2]Rivetline Concept Limited, Lagos, Nigeria
[3]Bisaj Engineering Limited, Osogbo, Nigeria
[*]Corresponding author: ridtiamiyu@gmail.com
D.O.I: 10.56201/ijgem.v8.no2.2022.pg144.172

Abstract
When a flood closes a rural road, the standard question is how much longer the trip takes. We argue that on sparse rural networks it is the wrong question, because the network usually does not offer a longer trip. Taking satellite-derived closure sets for the 2022 wet season over two Nigerian river basins and routing 10,438 settlement-pair trips across them, we find that between 47 and 60 percent of hit pair-dates have no path at all, and that among those that do reroute the penalty is small: a detour ratio of 1.003 to 1.100 and a worst realized delay of 4.8 hours against a median closure spell of 24 to 36 days. Absorption never fails on the margin. It fails by severance, and severance is local: under every closure input but the widest, every severed pair is a settlement cut into a pocket holding under 5 percent of network length, not a broken corridor. The links responsible are very few, eight out of 8,334 under the high-precision input. Assessed structurally instead, with no flood in the loop, the same network turns on three trunk links that sever 812 pairs each and that this detector puts under water on two acquisitions out of fifteen, under one closure input out of four, so a hazard-conditional ranking misses the largest exposure in the study area. Severed pairs are predictable in advance from structure alone, at an out-of-fold area under the curve of 0.9997, and the count of cut edges on the baseline path reaches 0.992 unaided, while betweenness centrality is the weakest of six predictors. But the identity of the critical links does not survive the detector: the top eight links under the high-precision and bounding readings of the same imagery share not one member. Structural fragility is robust; which fragile pair a given flood actually cut is a statement about the sensor. A depth-disruption test removes severance entirely under the high-precision input, the one that produces the headline, at all three depths swept, halves it under the bounding input at 300 mm and raises the count of merely delayed pairs by up to a factor of 2.5, so the split between severance and delay moves with the depth mapping as much as with the detector threshold.

Keywords: *network vulnerability, flood disruption, rural accessibility, rerouting, redundancy, remote sensing, Nigeria*

1 Introduction

A flood closes a road. The question the disruption literature asks is what the closure costs, and the currency it answers in is time: the additional travel time imposed on trips that must now go around (Jenelius et al. 2006; Jenelius and Mattsson 2012; Balijepalli and Oppong 2014). The apparatus is well developed and was built on dense networks, where going around is what happens. On a sparse rural network, it frequently is not. When roughly a quarter of the links are cut edges, as they are in both areas studied here, 26.9 percent in Kwara and 22.5 percent in Osun, a closure often does not lengthen a trip; it ends it. The costs then are not delay costs and are not measured in minutes, and an analysis that reports a mean travel-time penalty over the trips a closure touches will report a small number while a substantial share of those trips have no finite penalty to average. This paper takes that seriously and reorganizes the analysis around it. Using closure sets derived from Sentinel-1 imagery over the 2022 wet season for the Niger floodplain in Kwara State and the Osun River basin in Osun State, we route 10,438 settlement-pair trips over the observed network and take three questions in order. The first is whether the network absorbs the closure or severs the pair. The second is what kind of place fails where absorption fails, a corridor or a pocket. The third is whether those failures can be identified before the flood, from structure alone. A fourth question runs underneath all three. The closure sets are a sensor output with measured error, and that error is not one number. The accuracy assessment of the detector that produced them reports a precision of 0.008 on segments at watercourse crossings against 0.438 everywhere else, so precision is a statement about where in the landscape the detection falls; and at the operating point it selected its recall is 0.337, rising to 0.481 as the decision threshold is loosened to 0.05 and falling to 0.074 at the tightest backscatter threshold tested, so recall is a statement about how much the detector was asked to see. We therefore run every analysis under four readings of the same imagery, from high precision to a bounding envelope, and report how much of the spread in the answers is flooding and how much is the detector. The setting constrains what can be claimed. The networks are genuinely sparse: 26.9 percent of Kwara links and 22.5 percent of Osun links are cut edges, and only 41.5 percent of Kwara settlement pairs have an edge-disjoint alternative. The closure input carries a measured detection error, which lets the decision threshold be treated as a variable and swept rather than fixed at one value. And one of the two areas returns nothing: the detector found essentially no water in Osun, which we report as an absence of detection rather than an absence of flooding.

One scope condition is stated at the outset because it constrains the paper's central claim. The redundancy argument made here is a rural and corridor-scale argument, and it is not a claim about Nigeria's trunk network. Nothing measured in this paper bears on the trunk network either way. What is sparse here is the local and track network on which rural settlements actually depend. Section 2 places the work against the vulnerability and redundancy literatures. Section 3 describes the networks, the demand model and the four closure inputs. Section 4 gives the routing, absorption and predictability formalism. Section 5 reports the results and Section 6 takes up what they mean and what they do not support.

2 Related Work

2.1 Vulnerability, importance and exposure

Network vulnerability analysis asks what happens when a link is removed. Berdica (2002) set the terms, and the operational measures follow Jenelius et al. (2006), who separate a link's importance, the system-wide cost of losing it, from a location's exposure, the cost that location bears. Extensions cover area-covering disruptions (Jenelius and Mattsson 2012), robustness indices (Scott et al. 2006), critical segment identification (Sullivan et al. 2010; Bagloee et al. 2017), and the relationship between vulnerability and resilience (Mattsson and Jenelius 2015; Kermanshah and Derrible 2017). Balijepalli and Oppong (2014) give the network robustness measure closest in spirit to what we compute.

Nearly all of this work reports its results in travel time. That is the natural currency when the network reroutes, and on the urban and interurban networks these methods were developed for it almost always does. Jenelius (2010b) notes the distributional consequence, that peripheral users bear more, which is a step toward the present concern. What is less examined is the case where the reroute does not exist and the travel-time framing has nothing to measure.

2.2 Redundancy

The redundancy literature addresses this directly. Jenelius (2010a) treats redundancy as the property that makes importance finite. El-Rashidy and Grant-Muller (2016) and Xu et al. (2018, 2021) develop measures of how much alternative capacity a network holds, and Suurballe (1974) supplies the disjoint-path machinery we use to count alternatives. Percolation treatments (Li et al. 2015; Duan and Lu 2014) describe how connectivity collapses as links are progressively removed. What this literature establishes is that redundancy is the precondition for the standard analysis. What it has had less occasion to examine is a network where the precondition routinely fails, which is the rural African case and the one studied here.

2.3 Centrality as a predictor, and its limits

Betweenness centrality (Freeman 1977; Brandes 2001) is the usual first predictor of which links matter, supported by network efficiency (Latora and Marchiori 2001) and the primal treatment of street networks (Porta et al. 2006). Kirkley et al. (2018) show that in planar networks betweenness concentrates along a spanning structure. We report below that betweenness is the weakest of six predictors for the outcome that matters here, and the reason is structural: where cut edges are common rather than rare, a link's capacity to disconnect a pair is nearly independent of how much flow it carries.

2.4 Flood disruption specifically

Flood impacts on road networks have been modeled at multiple scales (Koks et al. 2019; Zhang and Alipour 2020; Papilloud and Keiler 2021), with attention to pluvial events (Yin et al. 2016; Pregnolato, Ford, Glenis, et al. 2017), preparedness (Arrighi et al. 2019) and emergency access (Coles et al. 2017; Green et al. 2017; Sohn 2006). Pregnolato, Ford, Wilkinson, et al. (2017) supply the depth-speed relation we use in Section 5.10 to test our binary closure rule. Observational studies of real disruptions (Zhu et al. 2010; Danczyk et al. 2017) document how travelers actually respond, against which our shortest-path reassignment is an idealization (Zhu and Levinson 2015; Lint et al. 2008). For Nigeria specifically, flooding on the Niger-Benue system (Nkeki et al. 2013), urban flooding and its modeling (Adelekan 2010; Nkwunonwo et al. 2020; Oyetubo et al. 2017; Echendu 2020) and rural accessibility (World Bank 2019; Ouma et al. 2018; Starkey 2016; Aderamo 2012) provide the setting. Population and settlement data follow Tatem (2017), network data follows OpenStreetMap with the completeness caveats of Barrington-Leigh and Millard-Ball (2017), and the regional network context is Blanford et al. (2012). The detection limitations under canopy that produce this paper's Osun null are documented by Makanga et al. (2017).

3 Study Areas, Network, and Closure Inputs

3.1 The two networks

The study covers the Niger floodplain in Kwara State and the Osun River basin in Osun State, the same two windows the satellite water detector covers. Settlement locations and populations are taken from the WorldPop 2020 constrained UN-adjusted layer and network geometry from

OpenStreetMap, both extracted for the two windows in December 2022, after the last acquisition in the series, so that a single network snapshot serves the whole season. Table 1 reports both networks.

Table 1. Network and demand inputs for the two study areas

	Detector segment table											
Area	Segs.	Nodes	km	Comp.[a]	% km in LCC[a]	Links	Comp.[b]	km	% km surf. tagged	Settl.	Population	OD pairs
Kwara	11,280	13,504	4,127	2,237	1.8	8,334	2	4,449	43.4	114	369,298	3,417
Osun	41,878	59,656	10,751	17,969	1.1	71,795	1	11,486	36.9	119	2,865,334	7,021

Note. The first three columns are the segment table produced by the satellite water detector applied to the road centerlines of these two windows. [a]Components and the share of network length in the largest component when the delivered node_from/node_to identifiers are used as given; the graph does not route, which is why this paper rebuilds topology from OpenStreetMap node identifiers held in the same cached Overpass response. [b]Components retained after the rebuild, each carrying at least 25 km; Kwara is two sub-networks because the Niger separates the banks and the nearest bridges lie outside the bounding box. Surface is the real OpenStreetMap surface or tracktype tag; the remainder is imputed at the measured class-conditional paved rate. Settlements are WorldPop 2020 constrained clusters of at least 500 persons, the largest 120 per area, snapped to the network. OD pairs are within-component pairs only.

Kwara carries 8,334 routable links over 4,449 km serving 114 settlements and 369,298 persons; Osun carries 71,795 links over 11,486 km serving 119 settlements and 2,865,334 persons. Kwara retains two components, of 2,606 and 1,843 km, split by the River Niger, which has no road bridge inside the window; the two banks are treated as separate networks because that is what they are. Osun retains one. Analysis runs on 3,417 within-component settlement pairs in Kwara and 7,021 in Osun. The structural facts matter more than the totals. 26.9 percent of Kwara links and 22.5 percent of Osun links are cut edges, whose removal disconnects the graph. Only 41.5 percent of Kwara pairs and 63.6 percent of Osun pairs have an edge-disjoint alternative path. Median baseline wet-season travel time is 170 minutes in Kwara and 96 in Osun.

3.2 Rebuilding a routable topology

The delivered segment table carries node_from and node_to identifiers, and a graph assembled from them cannot be routed on. Under those identifiers the 11,280 Kwara segments fall into 2,237 components, the largest holding 75 km out of 4,127. Identifiers that partition a network that finely are perfectly serviceable for the question the detector answers, which is whether a given 400 m segment was under water on a given date, and they carry no information about which segment adjoins which. Every result in this paper depends on adjacency, so the identifiers were set aside and topology was reconstructed from the OpenStreetMap node identifier lists held in the same cached Overpass response, matching all 11,280 Kwara and 41,878 Osun segments back by chainage with none unmatched. The diagnosis and the reconstruction are recorded in full with this paper's replication materials.

This data-preparation step is reported because the failure never announces itself. The join runs to completion, every field arrives populated, and any pair of settlements that happens to fall inside one fragment returns a plausible finite travel time. What a reuser would see is a network on which most trips are already impossible before any flood is applied, and there is no diagnostic in the output that distinguishes that from a genuinely disconnected landscape. Checking component structure against total length takes one line and is the only thing standing between the published product and a set of results that would read as catastrophic flood impact.

The link table analyzed here is narrower than the widest network these two windows would yield, and deliberately so. It reproduces the detector's own road-class filter exactly, so that the segment table and this link table describe the same roads, and it then discards every component carrying less than 25 km, which routing requires and a link-scoring exercise would not. Admitting the wider set of OpenStreetMap classes, service and living-street ways among them, would give a larger network and different totals from those reported above; the filter is stated here so that the totals can be reproduced.

3.3 Structural criticality before any flood

Removing each cut edge in turn and counting the settlement pairs it disconnects (Hopcroft and Tarjan 1973) gives a picture of the network's fragility that involves no flood at all. In Kwara, 68 links out of 8,334, 0.82 percent, can sever at least one pair. In Osun the same exercise gives 27 links out of 71,795, 0.04 percent, and a worst link severing 1.7 percent of pairs. That asymmetry, a 21.7-fold difference in the share of links capable of severing anything, is the structural difference between the two areas and it anticipates most of what follows.

The composition of the Kwara 68 matters for the policy reading and cuts against the easy version of it. Fifty-four are local roads, eleven are tracks and three are trunk links, and it is the three-trunk links that sit at the top. Links 154, 155 and 387 sever the same 812 pairs, 23.8 percent of all of them, which places them in series on a single trunk corridor: removing any one of them divides the larger of Kwara's two sub-networks, 2,606 km, into pieces of 1,745 and 860 km, both far above the 5 percent of network length that Section 4 uses to define a pocket. That is a corridor split in the strict sense, and it is the largest structural exposure in either area. Below the three the distribution falls away sharply. The next worst are three local links, again in series, severing 455 pairs and dividing that same sub-network into 2,097 and 509 km; the median severing link accounts for 71 pairs. Two exposures of different kinds therefore coexist in one network. A short trunk corridor whose loss would divide the study area, and a long tail of local roads and tracks each of which isolates one settlement. Which of the two a flood analysis finds turns out to depend on the closure input. Under the high-precision, calibrated and operating-point inputs no trunk link is closed on any of the fifteen acquisitions and every severing link is a track, a local road or, at the operating point, a tertiary one. Under the bounding input trunk link 387 closes on two acquisitions and severs 1,624 pair-dates on its own, and the bounding input is also the only one that records any corridor split at all. Whether a flood analysis of this network reports settlement isolation or corridor loss therefore turns on how one link is read on two dates, and Section 6 returns to that.

3.4 Demand

Flows between settlements were modeled with a gravity specification, no origin-destination survey existing for either area,

$$T_{ij} = \frac{P_i P_j}{d_{ij}^{\beta}}, \quad (1)$$

with P the WorldPop settlement population (Tatem 2017), d_{ij} the baseline network travel time and β held fixed at 2.0 and swept over 1.0 to 3.0. Because the flows come from a gravity specification rather than from observed counts, every headline in this paper is reported at the level of settlement pairs, which requires only that a pair exists, and flow weighting appears only as a sensitivity.

3.5 Four readings of the same imagery

The closure sets, the fifteen acquisition dates, the segment table, the operating point and every detection-skill figure quoted in this paper come from a Sentinel-1 water detector applied to the road centerlines of these two windows, together with the accuracy assessment that accompanies it. The detector outputs, the accuracy tables and the segment table are archived with this paper's replication materials, so every number attributed to them can be checked there. That detector does not produce one closure set; it produces a threshold, and the threshold is a choice with measured consequences. Table 2 gives the four inputs used here. The *calibrated* input takes every detection whose calibrated closure probability reaches 0.50. The *high-precision* input takes those and then drops the ones at watercourse crossings, the stratum in which the accuracy assessment measures a precision of 0.008 against 0.438 everywhere else. The *bounding* input takes any non-zero water fraction, which is looser than any decision threshold that assessment itself tested and where the loosest it did test returns a precision of 0.293, so it is an envelope on detected water rather than an estimate of flooding. The *operating point* is the threshold the assessment selected.

Table 2. The four closure inputs and what each one closes

Closure input	Rule	Kwara Seg.-days	Peak links	Peak km	% net.	Osun Seg.-days	Peak links	Peak km	% net.
High precision	conf. ≥ 0.50, no crossing	706	206	115.7	2.60	0	0	0.0	0.00
Calibrated	conf. ≥ 0.50	808	225	124.4	2.80	0	0	0.0	0.00
Bounding	water frac. > 0	11,410	2,330	1,139.5	25.61	4	1	0.5	0.00
Operating point	water frac. ≥ 0.25	5,070	1,377	696.5	15.66	0	0	0.0	0.00

Note. The rules verbatim are confidence >= 0.50 and not crosses_waterway, confidence >= 0.50, water_fraction > 0 and water_fraction >= 0.25, where confidence is the detector's calibrated closure probability and water_fraction its raw detector score. Segment-days count segment-by-acquisition observations flagged closed by the rule, over 15 Sentinel-1 acquisitions between 8 June and 5 December 2022. Peak links and peak km are the worst single acquisition. % net. is peak closed length as a share of the routable network. A link is closed when at least 25 % of its length is covered by closed segments. The first three inputs are nested: high precision ⊆ calibrated ⊆ bounding. The operating point is carried as a sensitivity, not a headline. High precision drops detections at watercourse crossings, a spatial stratum in which the accuracy assessment measures precision of 0.008 against 0.438 everywhere else. The two Osun facts in this row are both true and are easily read as contradicting each other. At its own selected threshold the detector returns no detection anywhere in Osun on any of the fifteen acquisitions, which is why three of the four Osun columns are zero; the bounding rule takes any non-zero water fraction, below that threshold as well as above it, and finds four such segment-days out of 628,170. Those four closes a single link, on four separate acquisitions, and that link lies on none of the 7,021 Osun baseline shortest paths, so every downstream Osun result is zero either way. Tree cover over 76 % of the Osun window and built-up double bounce are the reasons the accuracy assessment gives. Dropping the crossing detections removes 60 link-dates across 26 distinct links and 9.7 km, none of which lies on any baseline shortest path in this matrix, which is why the high-precision and calibrated columns of Table 3 are identical.

A link closes when the closed share of its length reaches a threshold held fixed at 0.25 and swept over 0.10, 0.25, 0.50 and 0.75. Closure is binary, which is the coarsest assumption in the paper and is tested rather than defended in Section 5.10.

3.6 Model assumptions

Ten quantities are set by the analyst, and the list below is complete.

Six are swept, and Section 5.9 reports the sweeps: the gravity deterrence exponent β, held fixed at 2.0 over 1.0 to 3.0; the closed-length share that closes a link, at 0.25 over 0.10 to 0.75; the settlement count, at 120 over 80 to 160; the minimum settlement population, at 500 over 250 to 1,000; the uniform scale on the speed table, at 1.0 over 0.5 to 2.0; and the hours of travel foregone per waiting day in Equation 3, at 8 over 4, 8 and 12. The depth mapping of Section 5.10 is a seventh, swept over the three depths its own source names.

Three are held fixed. The pocket threshold, at 5 percent of network length, decides whether a severance is recorded as a settlement cut off or as a corridor split; it governs how a severance is labeled and not how many are recorded. The minimum component length, at 25 km, decides which fragments of the rebuilt network are carried into the analysis at all. The maximum settlement snapping distance, at 5 km, decides which settlement centroids are admitted.

4 Methods

Two steps in the analysis are resampled and no others: the imputation of a paved or unpaved surface onto the links OpenStreetMap does not tag, which sets speeds, and the source sample used to approximate edge betweenness, which is one of the six predictors in Section 4.4. Each was repeated over eight independent draws, and the spread across them is reported in Section 5.7. Neither step can alter graph structure, a point that matters there.

4.1 Routing and the penalty

Let $G = (V, E)$ be the network, c_e the free-flow traversal time of link e, and $S \subset E$ the set of links closed on acquisition date t. For a settlement pair (i, j) write π_{ij} for the baseline shortest path and τ_{ij} for its travel time. The post-closure travel time is

$$\tau_{ij}^{t} = \min_{p \in \Pi_{ij}(G \setminus S_t)} \sum_{e \in p} c_e, \tau_{ij}^{t} = \infty \ \ if \ \Pi_{ij}(G \setminus S_t) = \emptyset, \quad (2)$$

and the penalty is $\delta_{ij}^{t} = \tau_{ij}^{t} - \tau_{ij}$. Three terms are used throughout and they nest. A pair-date is *hit* if $\pi_{ij} \cap S_t \neq \emptyset$, that is if a closed link lies on its baseline path. A hit pair-date is *severed* if $\tau_{ij}^{t} = \infty$ and *rerouted* otherwise. Hit is the union of the other two, so the severed share can never exceed the hit share, and every share reported in this paper uses these three words in these senses. We avoid the word *affected*, which in the disruption literature is used for both the hit set and the rerouted set and cannot carry both.

The distinction in Equation 2 between a large δ and an infinite one is the paper's organizing device, and it is the distinction a mean penalty destroys. We therefore report percentiles over all pair-dates, never a mean, and report the severed share separately from the rerouted penalty rather than folding one into the other. The detour ratio τ_{ij}^{t}/τ_{ij} is reported only over rerouted pair-dates, where it is finite by construction.

4.2 Absorption

Absorption asks whether the delay a closure imposes is bearable against the time the closure persists. Define the absorption ratio

$$A_{ij}^{t} = \frac{\delta_{ij}^{t}}{h\,\Lambda_t}, \quad (3)$$

with Λ_t the duration in days of the closure spell containing t and h the hours of travel a traveler forgoes on each day of waiting. Both numerator and denominator are therefore hours, which they would not be if the denominator were the calendar length of the spell: a spell is not a period of continuous travel, and dividing a detour by the wall-clock duration of a flood compares a journey with a season. h is assumed to be 8 hours, a working day, and is swept over 4, 8 and 12; since it enters linearly, halving it doubles every ratio reported below. Absorption *fails on the margin* when $A_{ij}^{t} \geq 1$, that is when going around costs more than waiting. Absorption *fails by severance* when $\delta_{ij}^{t} = \infty$. The two failure modes are conceptually distinct and, as Section 5 reports, only one of them occurs here. Taking the whole spell rather than a shorter tolerance is deliberate and works against the paper's argument. A traveler facing a closure does not necessarily wait for the water to recede; they may wait a day, or abandon the trip. Setting Λ_t to the full spell makes the denominator as large as it can plausibly be and marginal failure correspondingly as easy to observe, so a finding that absorption never fails on the margin under that denominator is stronger than the same finding under a shorter tolerance. We report a fixed six-hour comparison alongside for that reason.

Retained efficiency τ_{ij}/τ_{ij}^{t} is reported alongside, taking the value zero under severance.

4.3 Where absorption fails

Each severance is classified by what was cut. A *pocket* is a component of $G \setminus S_t$ holding under 5 percent of network length; a severance is an origin pocket, a destination pocket, or both, according to which endpoint falls in one. A *corridor split* is a severance in which neither endpoint is in a pocket, meaning the network has been divided into two substantial pieces. The classification separates local isolation from structural failure, and the two carry different policy implications.

4.4 Ex ante predictability

The operational question is whether the pairs that will be severed can be identified before the flood. We predict, for each pair, whether it loses connection on at least one acquisition, using six purely structural predictors that use no flood information: the count of cut edges on the baseline path, the maximum sampled edge betweenness along it, whether the pair has no edge-disjoint alternative, the count of articulation points on the path, the path length in links, and the baseline travel time. Validation is out of fold on five spatial blocks: a fold trains on pairs with neither endpoint in the held-out block and tests on pairs with both endpoints in it, so that no training pair shares a settlement with a test pair.

That design has a consequence which must be stated because it governs how the discrimination results below may be read. Requiring both endpoints in the held-out block leaves 1,199 of the 3,417 Kwara pairs ever scored out of fold. Every model quantity therefore lives on that subset and carries its own base rate, while the unfitted rules, which need no fold, are computed on all 3,417. The two are not comparable row for row and the table separates them.

Two further tests are reported alongside the headline discrimination. Univariate area under the curve for each predictor alone establishes whether the model is needed at all. And cross-input transfer, in which a model fitted on one closure input is tested against another input's target, asks whether what is being learned is the network or the detector. Transfer is computed twice: once on a full-data fit scored on full data, whose diagonal is by construction the in-sample figure, and once under the same spatial blocking as everything else, whose diagonal reproduces the out-of-fold figure exactly. Only the second is quoted next to an out-of-fold headline, and Section 6 states the further reason why even that number is weaker evidence than it looks.

4.5 Variance is reported in two parts and never pooled

Draw variance is the spread across the eight independent draws of the two resampled steps, the surface imputation and the betweenness source sample. *Area variance* is the difference between Kwara and Osun. These are reported separately and never added, because they answer different questions and because neither resampled step can in principle move a topological outcome: whether a pair is severed depends on graph structure, not on the speed assigned to a link. Section 5.7 reports both and the ratio between them. Where a ratio is reported it divides the area standard deviation by the draw standard deviation of Kwara, the area that carries the signal, and not by a draw standard deviation pooled across the two areas, which would be the same pooling this section rules out.

4.6 A test of the binary closure rule

Open-or-closed is the coarsest assumption made here. Rather than defend it, we test it against Pregnolato, Ford, Wilkinson, et al. (2017), who fit a limit vehicle speed to standing water depth and place impassability at 300 mm. Applying that function requires a depth, which the closure input does not carry, so link depth is set to D_{max} times the closed share of the link's length, a mapping that is assumed and is swept over $D_{max} \in \{150, 300, 450\}$ mm. The comparison is reported in Section 5.10 and it materially qualifies the paper's headline.

5 Results

5.1 The penalty distribution, and why no mean appears

Table 3 reports the Kwara penalty distribution over all 51,255 pair-dates under each closure input, and Figure 1 shows the distributions.

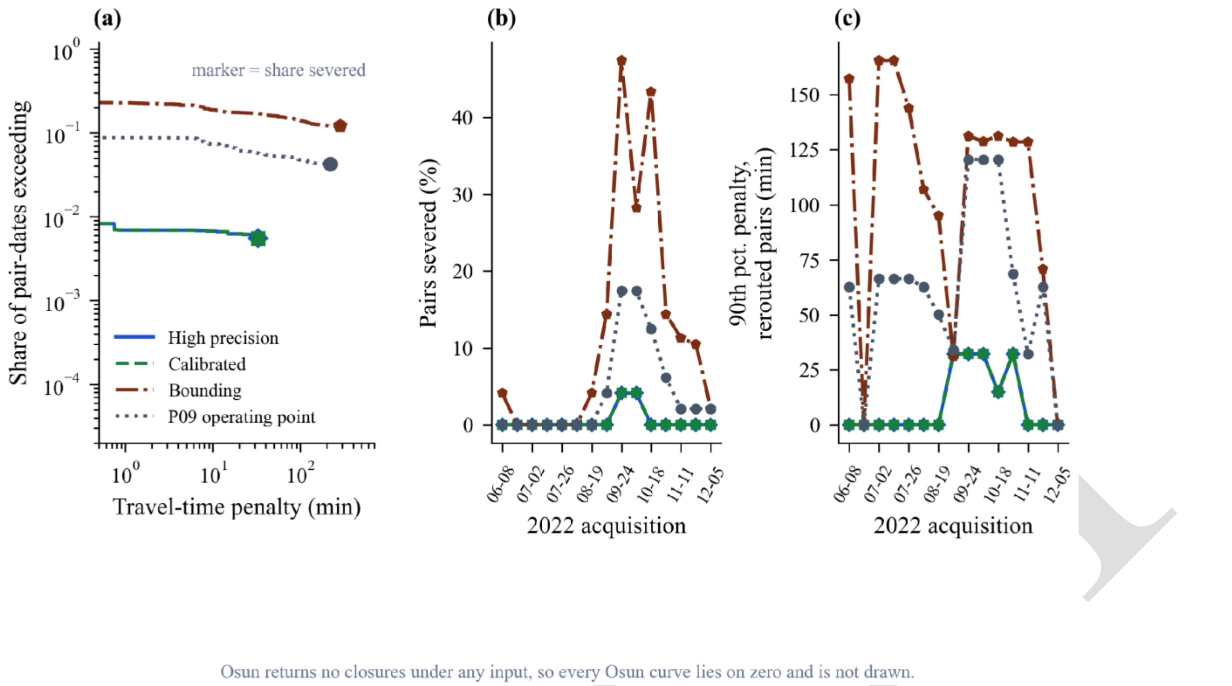

Figure 1. The Kwara penalty distribution under each closure input, over all 51,255 pair-dates. (a) Exceedance curve of the travel-time penalty on log-log axes, over all pair-dates rather than over the rerouted ones, so each curve terminates at the height of its own severed share; that terminal height is marked. Severed pair-dates have no finite penalty and so cannot sit on the curve, but omitting them is what produces a reassuring mean, and here they are the mass at which each curve stops. (b) Share of pairs severed on each of the fifteen 2022 acquisitions. (c) Ninetieth-percentile penalty over rerouted pairs on each acquisition. No Osun closure falls on a baseline shortest path under any input, so every Osun curve lies on zero and is not drawn.

Table 3. Distribution of the travel-time penalty over all origin-destination pair-dates

Closure input	% hit	% rerouted	% severed	Median	p90	p95	p99	Max	Median[a]	Detour ratio[a]
Kwara, Niger floodplain										
High precision	0.915	0.365	0.550	0.00	0.00	0.00	0.00	33.0	0.75	1.003
Calibrated	0.915	0.365	0.550	0.00	0.00	0.00	0.00	33.0	0.75	1.003
Bounding	23.287	11.291	11.997	0.00	6.81	40.75	131.32	286.5	15.16	1.100
Operating point	9.051	4.794	4.257	0.00	0.00	0.08	62.40	218.2	17.02	1.078
Osun, Osun River basin										
High precision	0.000	0.000	0.000	0.00	0.00	0.00	0.00	0.0	–	–
Calibrated	0.000	0.000	0.000	0.00	0.00	0.00	0.00	0.0	–	–
Bounding	0.000	0.000	0.000	0.00	0.00	0.00	0.00	0.0	–	–

Operating point	0.000	0.000	0.000	0.00	0.00	0.00	0.00	0.0	–	–

Note. Percentiles are minutes of additional wet-season travel time and are taken over all pair-dates, including the unaffected ones, which is why the median is zero everywhere; the mean is deliberately not reported, because a mean over a distribution with this much mass at zero and an infinite tail is not interpretable. The three shares use the definitions of Section 4: a pair-date is *hit* when a closed link lies on its baseline shortest path, *severed* when no path remains, and *rerouted* otherwise. Hit is the sum of rerouted and severed, so the severed share never exceeds the hit share. A severed pair-date has an infinite penalty and enters no percentile. [a]Median penalty and median ratio of post-closure to baseline travel time, over rerouted pair-dates only. Baseline is the wet-season network with nothing closed, so the quantity reported is a rerouting penalty and not a seasonal slowdown. Every Osun row is zero because the detector returned four segment-days of any water there across the whole season, closing one link that carries no baseline shortest path; those zeros are an absence of detection and must not be read as an absence of flooding, and the same caveat attaches wherever Osun is reported as zero or omitted in Tables 4 and 5.

Under the high-precision input a closed link lies on the baseline path of 0.915 percent of pair-dates. Of those, 0.550 percentage points are severed and 0.365 rerouted, so three-fifths of the pair-dates the flood touches lose their connection outright. Under the bounding input the three figures are 23.287, 11.997 and 11.291 percent, and at the operating point 9.051, 4.257 and 4.794. The severed share is a majority of the hit share under the two tightest inputs and close to half under the widest, which is the paper's central quantity and is stated as a share of what the flood touched rather than as a share of the whole matrix.

Where a reroute exists it is cheap. The detour ratio over rerouted pair-dates runs from 1.003 under the high-precision input to 1.100 under the bounding one, and the median penalty over rerouted pair-dates is 0.75 minutes and 15.16 minutes respectively. The worst realized delay anywhere in the study is 286.5 minutes.

On the worst acquisition of the season, 24 September 2022, at the operating point, 35.6 percent of pairs have a closure on their baseline path, 17.4 percent are severed, and the 18.2 percent that reroute face a ninetieth-percentile penalty of 120.6 minutes with a worst case of 218 minutes. Figure 2 maps the closures against the two networks on that date, and Figure 3 shows what four of the surviving reroutes look like on the ground.

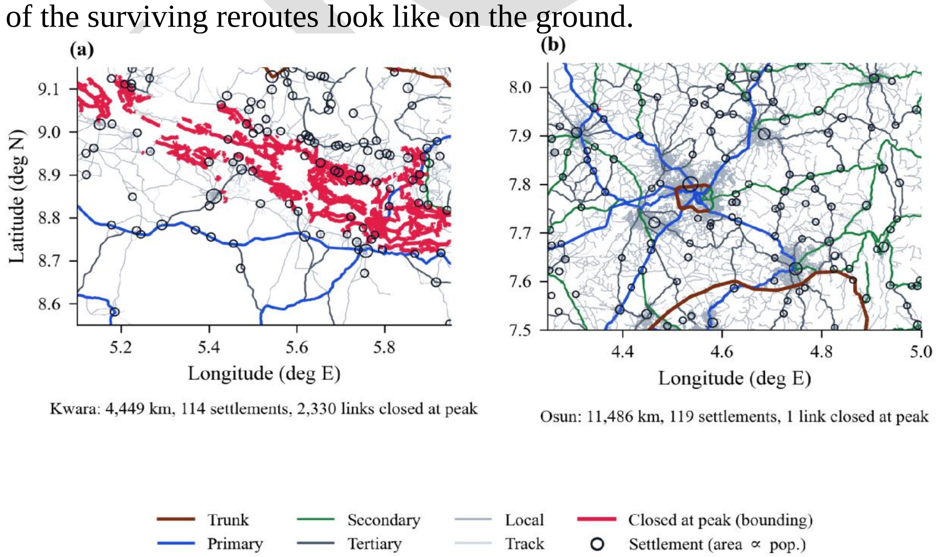

Figure 2. The two study areas with detected closures on 24 September 2022, the worst

acquisition of the season. (a) Kwara, the Niger floodplain; (b) Osun, the Osun River basin. Each panel shows the routable network by road class, the settlements with area proportional to population, and the links closed at the seasonal peak. Closures are drawn under the *bounding* input, the widest of the four, because it is the only one that puts enough length on the map to be legible at this scale: 2,330 closed links in Kwara against the 1,377 of the operating point and the 206 of the high-precision input. Osun carries four segment-days of any detected water across the whole season, closing one link, which is why every Osun figure in this paper is empty rather than small.

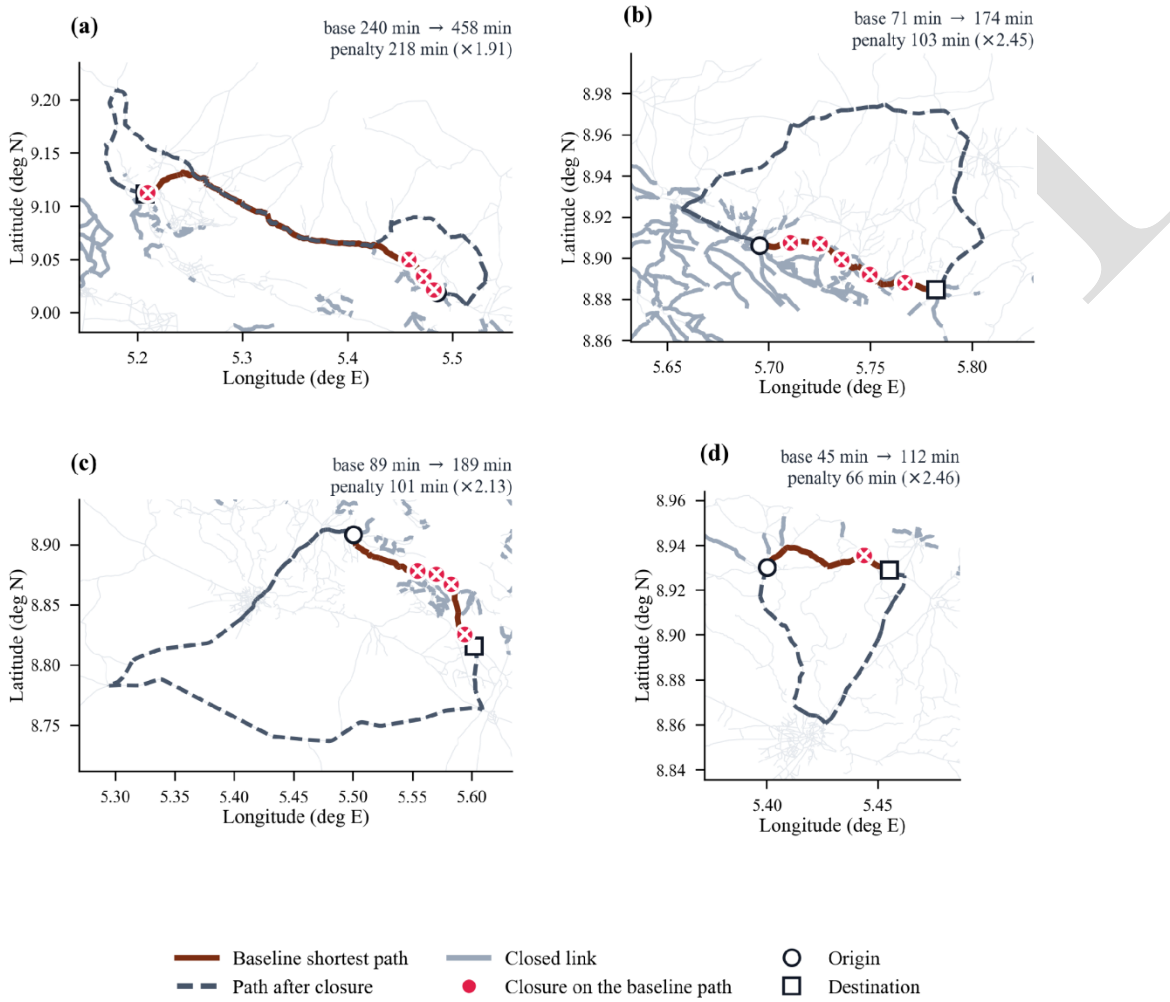

Figure 3. What a reroute looks like on the ground. Four Kwara settlement pairs at the detector's operating point on 24 September 2022, the peak acquisition, each mapped in longitude and latitude. The baseline shortest path is drawn solid, the path taken after the closure dashed, and the closed links that fall on the baseline path are marked; the surrounding network is drawn faint, with closed links picked out. The four are the largest surviving penalties on that date that share no closed link and no neighborhood with each other, so they are four distinct disruptions rather than four views of one corridor. Baseline and post-closure travel times are given on each panel. These are the visible tail of the penalty distribution, at ratios of 1.91 to 2.46; the median rerouted pair-date pays 1.078 at this operating point, which is a detour too small to see at this scale.

Two features of the table deserve comment. The ninetieth percentile is zero under three of the four inputs, because the overwhelming majority of pair-dates are untouched, which is why percentiles are reported over all pair-dates and no mean is reported anywhere. And Osun is zero on every row of every table: four segment-days out of 628,170 carry any non-zero water fraction, they close a single link, and that link carries no baseline shortest path, so no Osun pair-date is hit under any input.

5.2 Two closure inputs that differ by 13 percent and agree exactly

The high-precision and calibrated inputs give bit-for-bit identical results throughout, and the reason is informative about what detection error does.

Dropping watercourse-crossing detections removes 60 link-dates across 26 distinct links, 9.7 km, all of them local roads or tracks and 17 of them cut edges. Not one of those 26 links lies on any of the 3,417 baseline shortest paths. That is less surprising than it sounds once the base rate is stated: only 1,680 of the 8,334 Kwara links, 20.2 percent, carrying 1,220 of 4,449 km, lie on any baseline shortest path at all, so four links in five are invisible to this origin-destination matrix before any flood is applied. The false positives the accuracy assessment measures are real and numerous, but they sit on stretches of network carrying no settlement-to-settlement trip here, so removing 13 percent of the closed links changes the travel-time result not at all.

That is a statement about this origin-destination set, not a general result. A denser matrix, or one including intra-settlement trips, would touch those links. What it illustrates is that the impact of detection error depends entirely on whether the error falls where the demand is, which is not a property either the detector or its accuracy statistics can report.

5.3 Absorption fails by severance, never on the margin

Table 4 reports absorption and Figure 4 shows the failure structure.

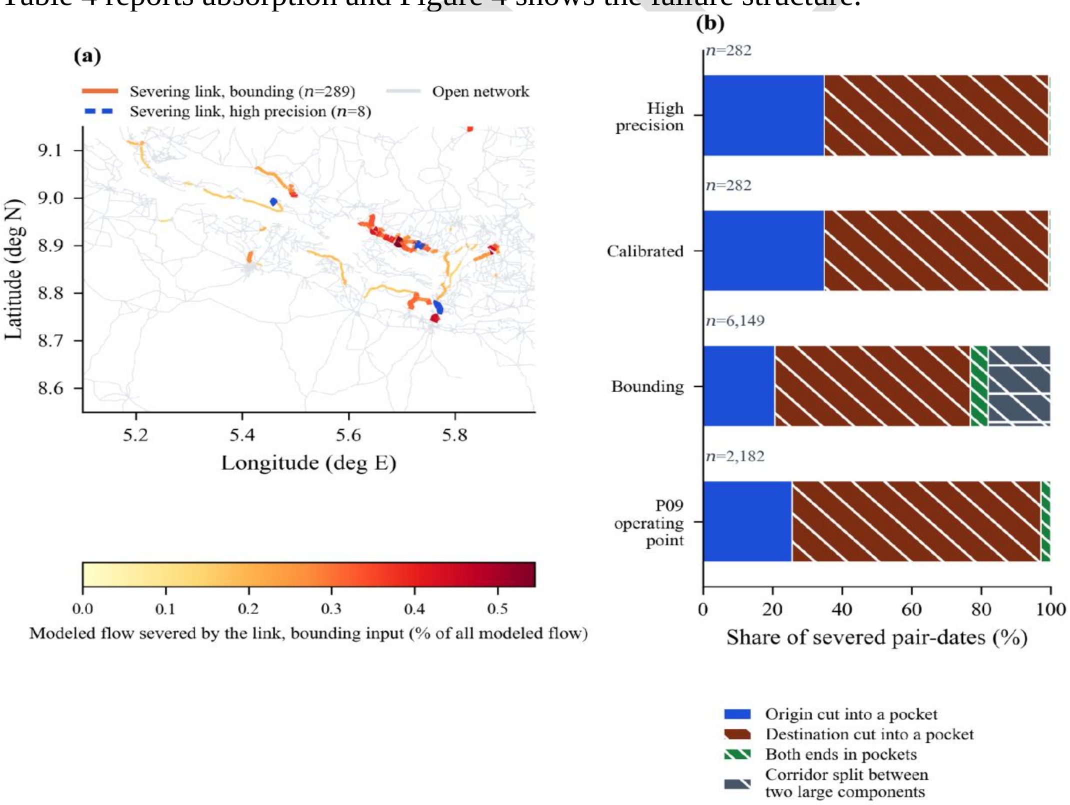

Figure 4. Where absorption fails in Kwara, and on what. (a) The links whose closure severs at least one settlement pair, drawn on the open network. Color and width give the modeled flow

each link severs under the bounding input, as a percentage of all modeled flow; the 289 such links under the bounding input are shown in that scale and the eight under the high-precision input are overdrawn in dashed blue. The two sets share no member. (b) Severed pair-dates decomposed by what was cut, as a percentage of each input's severed total, with that total printed above each bar. Destination pockets dominate throughout, and the corridor-split category is empty under every input but the bounding one.

Table 4. Absorption and where it fails, Kwara

Closure input	Hit	Rerouted	% severed	Wait (d)	Median R	Max R	% $R \geq 1$	% pocket	% corridor	Sev. links	% flow top 25
High precision	469	187	60.1	24.0	0.00013	0.0056	0.0	100.0	0.0	8	100.0
Calibrated	469	187	60.1	24.0	0.00013	0.0056	0.0	100.0	0.0	8	100.0
Bounding	11,936	5,787	51.5	36.0	0.00138	0.0497	0.0	81.9	18.1	289	56.8
Operating point	4,639	2,457	47.0	24.0	0.00139	0.0334	0.0	100.0	0.0	120	58.1

Note. A pair-date is *hit* when a closed link lies on its baseline shortest path, and hit is the sum of rerouted and severed. R is the absorption ratio of Equation 3: the realized penalty divided by the penalty a traveler with no alternative would bear, which is h hours of travel foregone on each day of the wait for the water to recede. h is assumed to be 8 hours and swept over 4, 8 and 12; halving it doubles every ratio and the count of marginal failures stays at zero throughout. Wait (d) is the median days to the first acquisition on which every closed link on the path is open again; the acquisition series is bounded by a twelve-day revisit and a twenty-four-day gap between 19 August and 12 September 2022, so this is a coarse quantity and a wait that never clears within the season is censored at one revisit past the last frame. $R \geq 1$ marks a rerouting that costs more than waiting; it never occurs, and the same holds against a fixed 6, 12 or 24 hour wait. Absorption therefore fails only by severance. % pocket and % corridor decompose severed pair-dates into those where an endpoint is cut into a component holding under 5 % of network length, and those where both endpoints remain in large components that are no longer joined. The 5 % pocket threshold is held fixed and governs how a severance is labeled, not how many are recorded. Sev. links is the number of distinct links whose closure severs at least one pair-date; % flow top 25 is the share of severed modeled flow carried by the 25 worst of them. Osun is omitted because no pair-date is hit there under any input.

Absorption never fails on the margin, and this is not a near miss. The worst absorption ratio anywhere in the study is 0.0497, and it falls on the same pair-date as the worst realized delay: a detour of 286.5 minutes, 4.8 hours, on a pair whose closure spell runs 12 days, over which a traveler choosing to wait instead forgoes 96 hours of travel at the assumed 8 hours a day. The median spell is longer still, 24 to 36 days depending on the input, which only makes waiting more expensive and rerouting more attractive. Even against a fixed six-hour tolerance rather than a spell, the worst ratio is 0.796 and nothing reaches one, and halving the hours foregone per waiting day merely doubles every ratio, which still leaves the count of marginal failures at zero. A traveler who can reroute always should, under any reasonable valuation of time.

Absorption fails by severance, and it fails often. Between 47.0 and 60.1 percent of hit pair-dates are severed depending on the input. Retained efficiency on hit pair-dates has a median of exactly zero under three of the four inputs, because more than half of hit pair-dates have no path at all.

This is the result that motivates the paper. A conventional analysis reporting mean delay over the trips a closure touches would have found a few minutes and concluded that these networks absorb flooding well. What they do is fail completely for half those trips and barely inconvenience the rest.

5.4 The failure is local isolation, not corridor loss

Decomposing severed pair-dates by what was cut, under every input except the bounding one, every single severance is a settlement cut into a pocket holding under 5 percent of network length. The corridor-split share is 0.0 percent under the high-precision, calibrated and operating-point inputs and 18.1 percent under the bounding one. Destination pockets dominate origin pockets throughout, between 56.2 and 71.6 percent against 20.7 to 34.8 percent.

The links responsible are very few. Under the high-precision input, eight links out of 8,334 ever sever a pair, 0.10 percent of the network. Under the bounding input 289 do, 3.5 percent, and the worst 25 of them carry 56.8 percent of all severed modeled flow.

The policy reading is direct but must be stated with the scope the previous section established. The exposure this flood reaches is a small number of single points of access to specific settlements, cheap to fix and easy to miss because nothing about their traffic volume marks them out. It is not the exposure the structural exercise found, which is a trunk corridor whose loss would divide the study area and which the detector places under water on only two acquisitions, and only under its widest reading. The two coexist and the flood analysis alone would not have found the second.

5.5 Severed pairs are predictable, and betweenness is the worst predictor

Table 5 reports out-of-fold discrimination and Figure 5 the curves.

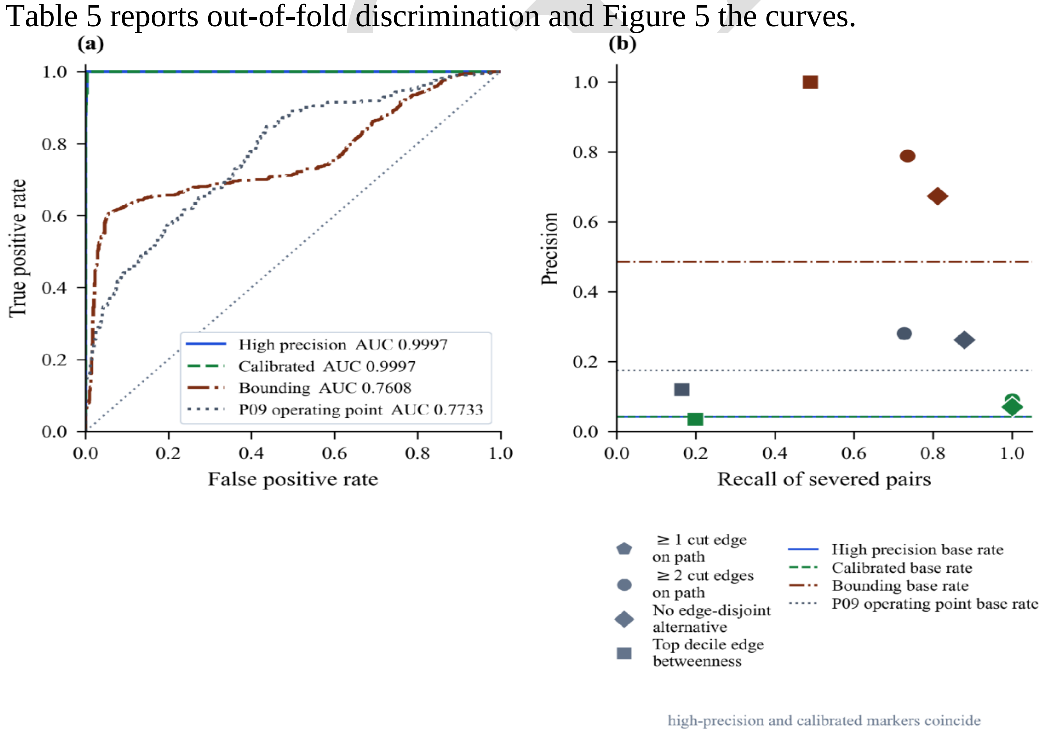

Figure 5. Identifying the severed pairs from structure alone. (a) Out-of-fold receiver operating

characteristic for the six-predictor severance model under each closure input, pooled across the five spatial blocks, with the area under each curve given in the legend. The high-precision and calibrated curves coincide. (b) Precision against recall for the four unfitted rules, computed on all 3,417 pairs, with each input's base rate drawn as a horizontal line so that distance above the line is the lift a rule buys. Marker shape gives the rule and color gives the closure input; the high-precision and calibrated markers coincide.

Table 5. Identifying the severed pairs before the flood, from network structure alone

Closure input / rule	n eval.	Base %	Pos.	AUC (o.o.f.)	AUC (in samp.)	Short-list	Prec.	Rec.	Lift
Fitted model, scored out of fold, shortlist the top 10 % of evaluated pairs									
High precision	1,199	2.75	33	0.9997	0.9981	120	0.275	1.000	9.99
Calibrated	1,199	2.75	33	0.9997	0.9981	120	0.275	1.000	9.99
Bounding	1,199	42.54	510	0.7608	0.8508	120	0.908	0.214	2.13
Operating point	1,199	17.60	211	0.7733	0.8425	120	0.617	0.351	3.50
Unfitted rule, applied to all 3,417 pairs, shortlist whatever the rule flags									
≥1 cut edge, high prec.	3,417	4.13	141	–	–	1,999	0.070	1.000	1.71
≥1 cut edge, calibrated	3,417	4.13	141	–	–	1,999	0.070	1.000	1.71
≥1 cut edge, bounding	3,417	48.52	1,658	–	–	1,999	0.673	0.812	1.39
≥1 cut edge, operating pt.	3,417	17.44	596	–	–	1,999	0.262	0.879	1.50

Note. The target is whether an origin-destination pair loses its connection on at least one of the fifteen acquisitions. Predictors are computed on the dry network and use no imagery, no rainfall and no closure information: cut edges on the baseline path, maximum sampled edge betweenness on that path, whether the pair has an edge-disjoint alternative, articulation points on the path, path length in links, and baseline travel time. Validation is spatial: settlements are partitioned into five k-means blocks, and a fold tests on pairs with both endpoints inside the held-out block while training on pairs with neither endpoint in it, so no training pair shares a settlement with a test pair. The two panels are computed on different populations and are not comparable row for row. Requiring both endpoints inside the held-out block leaves 1,199 of the 3,417 pairs ever scored out of fold, so every model quantity, and the base rate beside it, refers to that subset; the unfitted rule needs no fold and is applied to all 3,417. That difference in population, not a difference in skill, is why the out-of-fold AUC under the high-precision input exceeds the in-sample AUC computed on all pairs. Short-list is the number of pairs shortlisted, and precision, recall and lift are computed on that shortlist. The two panels shortlist on different rules and at very different sizes, so the three columns are not comparable across the midrule. For the fitted model the shortlist is the highest-scoring 10 % of the 1,199 evaluated pairs, 120 pairs, and lift is against the base rate of that evaluated subset. For the unfitted rule the shortlist is every pair whose baseline path uses at least one cut edge, 1,999 pairs or 58.5 % of all 3,417, and lift is against the base rate over all pairs; a rule that flags 58.5 % of the matrix can and does reach a recall of 1.000 at a precision near the base rate. The count of cut edges on the baseline path, used alone as a score, reaches an area under the curve of 0.992 under the high-precision and calibrated inputs, 0.810 under the bounding input and 0.769 at the operating point. Osun cannot be assessed: it has no positives under any input, for the reason given under Table 3.

The severed pairs are identifiable in advance from structure alone. Out-of-fold area under the curve is 0.9997 under the high-precision and calibrated inputs, 0.7733 at the operating point and 0.7608 under the bounding input, where the target is far commoner and the problem correspondingly harder. Those figures come with a qualification about population that Table 5 carries in full. The spatial blocks are strict: a fold tests only on pairs with both endpoints inside the held-out block, which means 1,199 of the 3,417 Kwara pairs are ever scored out of fold, and the severance rate within that evaluated subset is 2.75 percent under the high-precision input against 4.13 percent over all pairs, and 42.5 percent under the bounding input against 48.5 percent. Every model number quoted in this section is a statement about those 1,199 pairs. It also explains an anomaly a reader would otherwise stop at, that out-of-fold discrimination under the high-precision input, 0.9997, is higher than in-sample discrimination, 0.9981: the two are computed on different sets of pairs, and the held-out blocks happen to contain an easier problem than the matrix as a whole, not a model that improves when it is denied training data. One predictor does nearly all of the work. The count of cut edges on the baseline shortest path, alone, reaches an area under the curve of 0.992 under the high-precision input and 0.769 at the operating point. The simplest usable rule, flagging every pair whose baseline path uses at least one cut edge, needs no fold at all and can be applied to the whole matrix: it flags 58.5 percent of all 3,417 Kwara pairs and catches 100 percent of the pairs severed under the high-precision and calibrated inputs, 87.9 percent at the operating point and 81.2 percent under the bounding input. No model is required to act on this, and unlike the model figures these are statements about every pair in the study area. The gap between the 0.9997 of the full model and the 0.992 of the single predictor is not worth the model, and the difference in coverage widens the gap further in the rule's favor. The rule is auditable by an engineer with a map, requires no fitting and therefore cannot be overfitted, and degrades gracefully: under the hardest input, where the base rate over all pairs is 48.52 percent and the fitted model falls to 0.7608 on the pairs it can be scored on, the rule still catches 81.2 percent of severances across the whole matrix. A screening instrument that a district engineer can apply without software is worth more here than three additional points of discrimination on a minority of the pairs. Betweenness centrality is the weakest of the six predictors, with a univariate area under the curve of 0.612 under the high-precision input. That reverses what the urban network literature would suggest (Freeman 1977; Kirkley et al. 2018) and the reason is structural rather than anomalous: where cut edges make up 26.9 percent of the network, a link's capacity to disconnect a pair is very nearly independent of how much flow crosses it. High betweenness identifies links that carry a lot; it does not identify links that are the only way through, and on a dense network those are usually the same links while here they are not. A broader reading is available, on two areas and one season. Centrality measures were developed on networks where flow and structure are coupled, and much of their predictive success may be borrowed from that coupling rather than intrinsic. Where the coupling breaks, as it does on a network in which 22.5 to 26.9 percent of links are cut edges, the measure that survives is the one that speaks directly to connectivity. Whether that holds beyond these two basins is a question this study can raise and not answer.

Osun cannot be assessed on this question. It has no positives under any input.

5.6 How much of the answer is the detector

Figure 6 reports the comparison across closure inputs on the same network, the same season and the same flood.

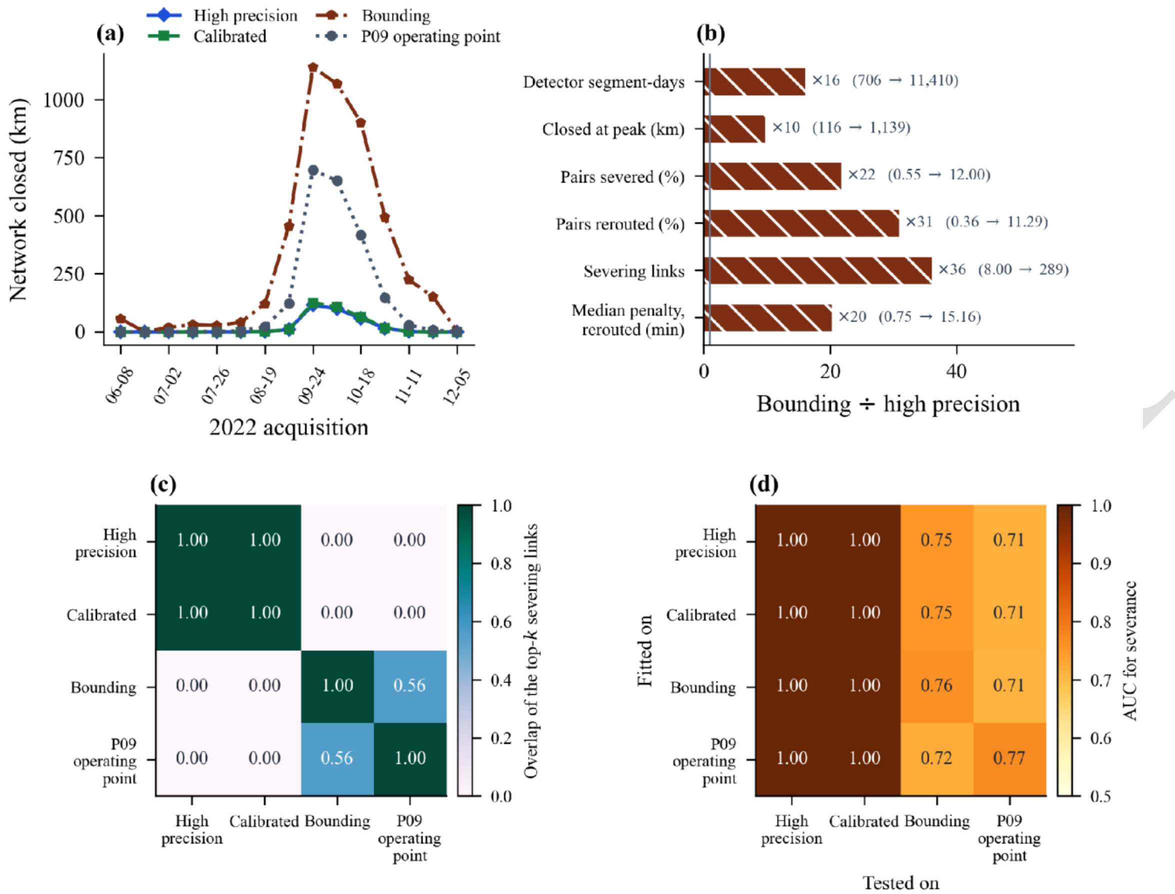

Figure 6. The same flood read four ways. (a) Kilometers of the Kwara network closed on each of the fifteen 2022 acquisitions, under each closure input. (b) The ratio of the bounding input to the high-precision input on six headline quantities, on a linear scale, with the underlying pair printed beside each bar and a vertical line at unity. (c) Overlap of the top-k severing links between each pair of inputs, where k is the smaller of the two sets: high precision and calibrated share 8 of 8, high precision and bounding share 0 of 8, bounding and the operating point share 67 of 120. (d) Area under the curve when a severance model fitted on the input named by the row is tested against the target of the input named by the column, under the same spatial blocking used elsewhere in the paper, so that the diagonal is the out-of-fold figure of Table 5.

The bounding input exceeds the high-precision input by a factor of 16 on detector segment-days, 10 on network length closed at peak, 22 on pair-dates severed, 31 on pair-dates rerouted, 36 on the count of links that ever sever a pair, and 20 on the median penalty over rerouted pair-dates. These are not competing estimates of different quantities; they are the same quantity computed from the same imagery at two thresholds.

The shortlists do not agree at all. Ranking links by severed modeled flow and comparing the top k, where k is the smaller of the two sets, the high-precision and calibrated inputs share 8 of 8. The high-precision input shares *0 of 8* with the bounding input and *0 of 8* with the operating point. The bounding input and operating point share 67 of 120.

Not one of the eight links that sever a pair under the high-precision reading appears anywhere in the bounding reading's top eight. If the object of the exercise is a list of links to fix, the detector threshold determines the list.

Something does transfer, and it needs to be stated with more care than the shortlist result. Under the same spatial blocking used everywhere else, a severance model fitted on the bounding input scores 1.000 against the high-precision target, and one fitted on high precision scores 0.752 against the bounding target. Read at face value, that says which pairs are structurally fragile is robust to the detector while which of them a given flood actually cut is not.

Two things weaken it. The first is that these numbers are asymmetric for a mechanical reason. The targets are strictly nested: every one of the 141 pairs severed under the high-precision input is also severed under the bounding input, and every one of the 596 severed at the operating point is too. A model fitted on the 1,658 bounding positives is therefore being asked, in the bounding-to-high-precision direction, to rank a 141-member subset of its own training target, and scoring 1.000 on that is close to arithmetic. The reverse direction, 0.752, is the informative one and it is the weaker number. The second is that the high-precision target has 33 positives among the 1,199 pairs ever scored out of fold, which is a thin basis for a discrimination statistic however it is computed. The transfer is reported because the asymmetry itself is informative, and Section 6 returns to it.

5.7 Imputation-draw variance and area variance

Table 6 reports the two variances separately for five headline quantities, and they are not of the same order.

Table 6. Imputation-draw variance and area variance, reported separately

Quantity	Kwara mean	Osun mean	Draw s.d. Kwara	Draw s.d. Osun	Area s.d.	Area/draw
High precision						
Pair-dates severed (%)	0.5502	0.0000	0.00000	0.00000	0.38904	–
Pair-dates rerouted (%)	0.3648	0.0000	0.00000	0.00000	0.25798	–
90th pct. penalty (min)	0.0000	0.0000	0.00000	0.00000	0.00000	–
99th pct. penalty (min)	0.0000	0.0000	0.00000	0.00000	0.00000	–
Median penalty, rerouted (min)	0.7464	0.0000	0.00000	0.00000	0.52778	–
Calibrated						
Pair-dates severed (%)	0.5502	0.0000	0.00000	0.00000	0.38904	–
Pair-dates rerouted (%)	0.3648	0.0000	0.00000	0.00000	0.25798	–
90th pct. penalty (min)	0.0000	0.0000	0.00000	0.00000	0.00000	–
99th pct. penalty (min)	0.0000	0.0000	0.00000	0.00000	0.00000	–
Median penalty, rerouted (min)	0.7464	0.0000	0.00000	0.00000	0.52778	–
Bounding						
Pair-dates severed (%)	11.9969	0.0000	0.00000	0.00000	8.48307	–
Pair-dates rerouted (%)	11.8562	0.0000	0.36063	0.00000	8.38357	×23.2
90th pct. penalty (min)	7.4984	0.0000	0.41713	0.00000	5.30220	×12.7
99th pct. penalty (min)	134.8118	0.0000	4.69174	0.00000	95.32632	×20.3
Median penalty, rerouted (min)	18.2505	0.0000	1.90963	0.00000	12.90503	×6.8

Note. Draw s.d. is the standard deviation across 8 independent imputation draws within an area; area s.d. is the standard deviation of the two area means. Kwara mean and Osun mean are means across those same 8 draws, so wherever the draw s.d. is non-zero they differ a little from the corresponding headline figure in Tables 3 and 4, which are computed on the single draw the rest of the paper uses: the rerouted share reads 11.86 % here against 11.291 % there, and the ninetieth-percentile penalty 7.50 against 6.81 minutes. The severed share is identical in the two places because its draw s.d. is exactly zero. The two are never pooled, and the last column divides the area s.d. by the Kwara draw s.d. printed three columns to its left, so a reader can reproduce it from the table. The only resampled steps in the analysis are the imputation of a paved or unpaved surface onto links OpenStreetMap does not tag and the source sample for the betweenness predictor; neither can move a topological outcome, so the share of pair-dates severed has a draw standard deviation of exactly zero and the ratio is undefined rather than large. Where the ratio is defined it runs from ×6.8 to ×23.2, which brackets the ×19 the accuracy assessment reports for the detector's own detection skill. Osun contributes a mean of zero to every row because it returns no closures. Sweeps over every assumed parameter are reported in full with the replication materials: closing a link at a 0.10 rather than 0.25 length share raises Kwara severance under the bounding input from 12.00 to 14.62 %, and at 0.75 lowers it to 7.77 %; changing the settlement count from 80 to 160 moves it from 9.89 to 12.78 %; halving the hours of travel foregone per waiting day doubles every absorption ratio and leaves the count of marginal failures at zero; a uniform scale on the speed table divides every absolute penalty by that scale and leaves every share and every rank unchanged. Replacing the binary open-or-closed rule with the depth-disruption function of Pregnolato and colleagues, under an assumed mapping from closed length share to depth, halves Kwara severance under the bounding input at their 300 mm impassability depth (11.997 to 6.977 %) while raising the share merely delayed from 11.291 to 21.530 %, and removes severance entirely at 150 mm; under the high-precision input severance is zero at all three depths. See Figure 7(d).

Under the bounding input, the share of pair-dates severed has a draw standard deviation of exactly zero, and the corresponding area standard deviation is 8.483 percentage points. The ratio is undefined rather than large, and the reason is structural: the surface imputation sets link speeds and nothing else, and no speed can change whether a path exists.

Where the ratio is defined it runs from 6.8 to 23.2, each figure being the area standard deviation divided by Kwara's own draw standard deviation, since Osun's is identically zero and pooling the two would be the operation this section exists to refuse. The share of pair-dates rerouted has a draw standard deviation of 0.361 percentage points against an area standard deviation of 8.384, a factor of 23.2. The ninetieth-percentile penalty gives 12.7, the ninety-ninth 20.3, and the median penalty over rerouted pair-dates 6.8. That range brackets the factor of 19 the accuracy assessment reports for the detector's own detection skill, which is a useful coincidence rather than a validation. The two variances are reported separately and are never added. Draw variance says the analysis is stable under resampling; it is, to the point of exactness on the topological outcomes. Area variance says Kwara and Osun are different places, and since Osun is identically zero throughout, the area variance is almost entirely the statement that the detector saw nothing there.

5.8 Which reading is closest to the flood is not settled here

Nothing above says which of the four inputs is nearest the truth, because the paper's design treats the spread between them as the object of study rather than as a nuisance to be resolved. Settling it would take an independent observation of flooded road length in these two windows in 2022, matched to this network and to these acquisition dates, and no such observation was available to us. The question is left open, which is why every result in this paper is reported under all four inputs rather than under a preferred one.

What the inputs themselves settle is narrower, and the paper acts on it throughout. The bounding rule takes any non-zero water fraction, which is looser than any decision threshold the detector's own accuracy assessment tested, and the loosest it did test returns a precision of 0.293. The bounding input is therefore an envelope on what a sensor detected and must not be read as an estimate of what happened. That is a statement about the rule, not a ranking of the four readings against the flood.

5.9 Sensitivity to every assumed parameter

Figure 7 reports the sweeps. Kwara severance under the bounding input moves from 14.62 to 7.77 percent as the closed-length threshold goes from 0.10 to 0.75, from 9.89 to 12.78 percent as the settlement count goes from 80 to 160, and from 12.00 to 10.41 percent as the minimum settlement population goes from 250 to 1,000. It is unchanged at 11.997 percent across a speed scale from 0.5 to 2.0, for the same reason the surface imputation cannot move it.

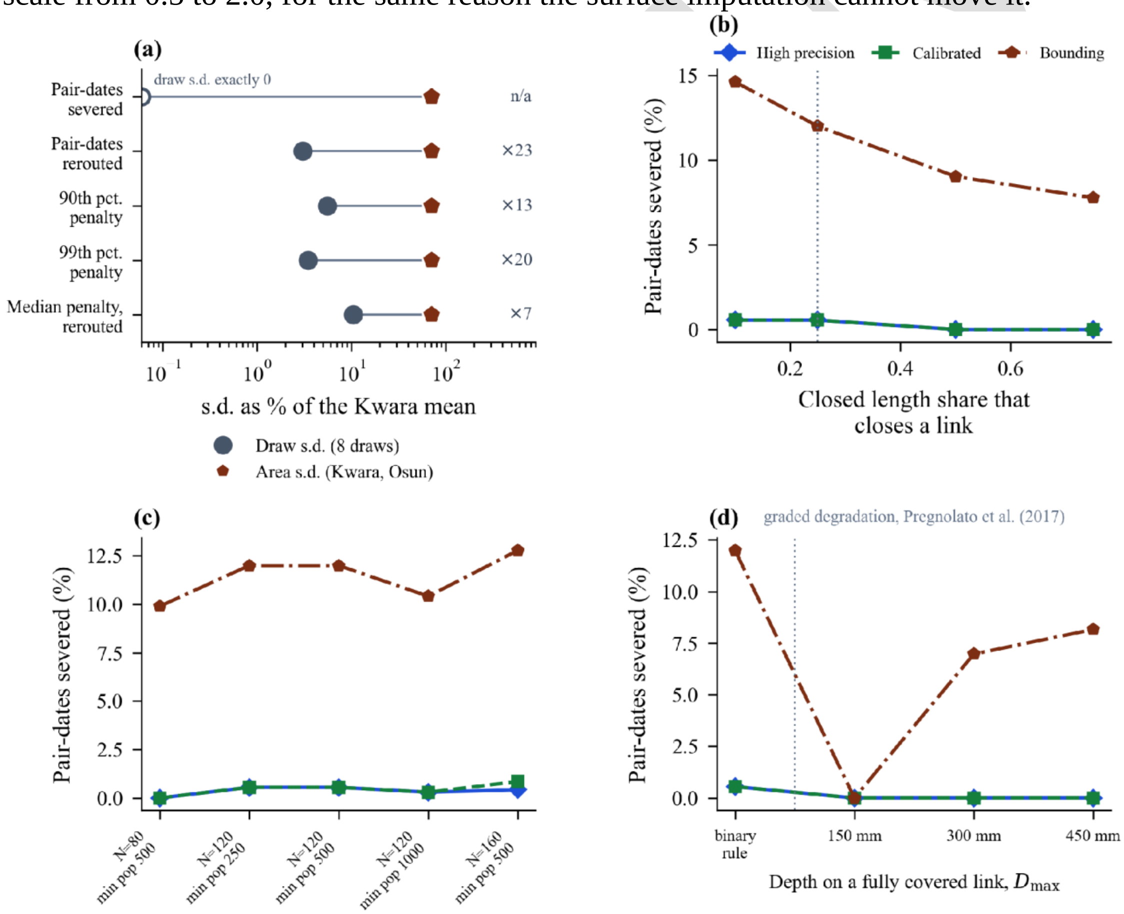

Figure 7. Variance and sensitivity. (a) Imputation-draw standard deviation against area standard deviation for five Kwara quantities under the bounding input, both as a percentage of the Kwara mean on a log axis, with the ratio between them printed at the right. The draw standard deviation of the severed share is exactly zero and is drawn as an open marker at the axis floor, because a topological outcome cannot move with a speed. (b) Kwara severance against the closed-length

share that closes a link, with the assumed 0.25 marked. (c) Kwara severance under each of the five settlement definitions swept. (d) Kwara severance under the binary rule against the graded depth-disruption function of Pregnolato, Ford, Wilkinson, et al. (2017) at each swept depth. Severance under the bounding input halves at 300 mm and reaches zero at 150 mm, and under the high-precision and calibrated inputs, whose curves coincide on the axis, it is zero at every depth.

The gravity exponent moves an answer further than any other assumed quantity, and further than the detector does. The severed flow share runs from 3.4789 percent at $\beta = 1.0$ to 0.0017 percent at $\beta = 3.0$, a factor of 2,091 across three orders of magnitude. It moves only the flow-weighted quantities and leaves every pair-level result untouched, which is precisely why the pair-level results are the headline and flow weighting appears as a sensitivity.

The sweep on the closed-length threshold runs in the direction one would expect and is worth reading carefully, because it moves severance the wrong way. Raising the threshold from 0.10 to 0.75 requires more of a link to be under water before it closes, which closes fewer links, and severance falls from 14.62 to 7.77 percent. That is a factor of 1.9 across the full plausible range of that threshold, against a factor of 22 between the tightest and widest readings of the detector. The threshold matters; the detector matters more.

The settlement sweep moves in the opposite direction, severance rising from 9.89 to 12.78 percent as the settlement count goes from 80 to 160. Adding settlements adds peripheral ones, and peripheral settlements are the ones that sever, which is a restatement of the pocket finding rather than an independent result.

5.10 The binary closure rule, tested

Applying the depth-disruption function of Pregnolato, Ford, Wilkinson, et al. (2017) under the assumed depth mapping changes the headline substantially, and the size of the change depends on which input is read. Each pair below is taken at one depth, because severance and delay must be compared at the same depth or the comparison means nothing.

Under the high-precision input, the one that produces the paper's 0.550 percent severance headline, its eight critical links and the whole of Section 6.4, the graded treatment returns severance of **0.000 percent at 150 mm, at 300 mm and at 450 mm alike**, while the share of pairs delayed rises from 0.365 percent under the binary rule to 0.915 percent at every one of the three depths. The calibrated input, identical throughout, gives the same. Under a published graded rule and this depth mapping, the high-precision input therefore produces no severance at any depth swept.

Under the bounding input, severance falls from 11.997 percent to 6.977 percent at $D_{max} =$ 300 mm, the depth those authors call impassable, while the share of pairs merely delayed rises from 11.291 to 21.530 percent at that same depth, a factor of 1.91. At 150 mm severance reaches 0.000 percent and delay 28.506 percent, a factor of 2.52. At the operating point severance falls from 4.257 to 1.996 percent at 300 mm while delay rises from 4.794 to 9.047 percent at 300 mm; the 11.043 percent that a reader might expect to see quoted here is the 150 mm figure, where severance is zero rather than halved. Across every input and every depth swept, the largest observed multiple on the count of pairs merely delayed is 2.52, so the change is a doubling and falls well short of a tripling.

The binary rule converts what a graded treatment would call widespread slowing into localized severance. The split between severance and delay is therefore as sensitive to the depth mapping as it is to the choice of closure input.

The geography is unchanged. Which pairs are hit and which links carry the flow are the same under the depth treatment as under the binary rule; what moves is the balance between delay and severance?

6 Discussion

6.1 The travel-time framing does not fit this network

The central finding is that between 47 and 60 percent of the pair-dates a flood touches have no path at all, while the ones that reroute do so at a detour ratio between 1.003 and 1.100. There is no middle of this distribution. A trip is either essentially unaffected or it is impossible, and the mean delay that a conventional vulnerability analysis reports would be a number describing neither group. This is not a criticism of the vulnerability literature, which developed its measures for networks that reroute and measures them well. It is an observation that the measures carry an assumption, redundancy, that the rural networks in this study do not satisfy, and that the assumption is invisible in the output. A mean penalty of a few minutes reads as resilience whether it arises because the network absorbed the closure or because the trips it could not absorb were dropped from the average. The practical replacement is not complicated. Report the severed share and the rerouted penalty as two numbers, never as one; state both as shares of the pair-dates the flood touched, not of the whole matrix; and report percentiles over all pairs rather than a mean over the touched ones.

6.2 Two exposures at two scales, and the analysis sees only one

The redundancy argument here is rural and corridor-scale, as stated at the outset. One version of this paper's finding is easy to write and is not supported by the data.

That version says the exposure is entirely capillary: that the links which can sever a settlement pair are the last tracks into villages, that a national trunk network can be redundant while every village hangs off one road, and that the two statements describe different networks and cannot conflict. The first clause is false in this study area. Of the 68 Kwara links whose removal severs a settlement pair, 54 are local roads and 11 are tracks, but three are trunk links, and those three are the most consequential links in the network. They sever 812 pairs each, 23.8 percent of the matrix, against 455 for the next group and 71 for the median severing link, and removing any of them divides the north bank into two pieces of 1,745 and 860 km. That is a corridor split by this paper's own definition, at the top of its own ranking, on a trunk route.

What is true, and is the more interesting claim, is that the flood analysis does not find it. No trunk link is closed on any acquisition under three of the four closure inputs; under the fourth, the widest, one trunk link closes on two of fifteen acquisitions. So the exposure that dominates the network structurally is almost invisible to the exposure measured from imagery, and the two exercises, run on the same network in the same season, return different objects. A flood-conditional analysis reports the links this flood happened to reach as read by this detector. A structural analysis reports the links that matter if anything reaches them. Neither is a substitute for the other, and a study that runs only the first will report a capillary problem in a network that also has a corridor problem. The methodological consequence generalizes past the present case. Vulnerability results are routinely reported without stating either the scale of the network analyzed or whether the criticality is structural or hazard-conditional, and those two axes give different answers here. At settlement scale and conditional on this flood, this network is a collection of single points of access. At corridor scale and unconditional, it turns on three trunk links. A vulnerability assessment that does not say which network it analyzed and whether the hazard was in the loop has not said enough for a reader to know what it found.

6.3 The detector threshold and the ranking of critical links

The sharpest methodological result is the shortlist comparison. The eight links that sever a pair under the high-precision reading share no member with the top eight under the bounding reading, on identical imagery. Any exercise whose output is a prioritized list of links to protect is, on this evidence, reporting its detector threshold as much as its network.

The transfer result appears to cut the other way, and the nesting of the targets accounts for most of it. Held out under the same spatial blocking as everything else, a severance model fitted on one closure input predicts another input's target at between 0.715 and 1.000, which reads as structural fragility being recoverable whichever imagery one adopts. The high transfer figures are a model ranking a subset of its own training target, as Section 5 sets out, and the direction that is not mechanical returns 0.752, which is useful and unspectacular.

The unfitted structural screen is the more useful instrument here: the set of pairs whose baseline path crosses a cut edge. It requires no fitting, no folds and no flood data, is computed on every pair rather than a held-out minority, and catches 100 percent of high-precision severances and 81.2 percent under the widest input. It makes no claim about a particular flood, and the case for it is independent of the transfer result.

6.4 Implications for practice

The cheapest instrument in this paper is also the most useful one. Flagging every pair whose baseline path uses at least one cut edge flags 58.5 percent of Kwara pairs and catches every high-precision severance, and that calculation needs a road network and a settlement layer, both freely available, and no imagery. What such a screen should look for is a single point of access, because that is what a flood of this kind actually cuts. Under every closure input but the widest, every severance recorded here is a settlement cut into a pocket, so the remedy is a second access track rather than a raised highway, and the candidate set under the high-precision input is eight links out of 8,334. An investment program should not stop at what the flood reached. The same network holds three trunk links whose loss would split the study area, and the flood analysis reaches them on two acquisitions under one input out of four, so a program that ranks only what a detector saw under water will not have those three links on any list. The structural exercise is worth running alongside it, and worth running without the hazard in the loop. Nor should a detection threshold be left to choose the investment list. Where a ranking is required, it should be reported under at least two readings of the imagery, with the disagreement between them treated as the honest error bar. Here that disagreement is total.

6.5 Limitations

Two limitations govern how these numbers may be used.

Osun. The detector returned essentially nothing there, so this analysis returns nothing there. The Osun River basin was not flood-free in 2022; the method did not see it, plausibly because of canopy (Makanga et al. 2017). Every Osun zero in this paper is an absence of detection and must not be read as an absence of flooding.

Rerouting and behavior. Rerouting here is shortest-path reassignment (Zhu and Levinson 2015) with no volumes, capacities or assignment, so every penalty is a free-flow lower bound on the network consequence of a closure. Real travelers cancel trips, wait, switch mode, and take routes a shortest-path model would not choose (Danczyk et al. 2017; Lint et al. 2008).

7 Conclusions

We routed 10,438 settlement-pair trips across satellite-derived flood closure sets for the 2022 wet season in two Nigerian river basins, under four readings of the same imagery, and asked whether the networks absorb closures or fail.

They fail, and they fail in a way the standard travel-time framing cannot express. Between 47 and 60 percent of the pair-dates a closure touches have no path at all, while pairs that do reroute face a detour ratio of 1.003 to 1.100 and a worst realized delay of 4.8 hours against a median closure spell of 24 to 36 days. Absorption never fails on the margin; it fails by severance. Under every closure input but the widest, every severance is a settlement cut into a pocket, and eight links out of 8,334 are responsible under the high-precision input.

The failures are predictable in advance from structure alone, at an out-of-fold area under the curve of 0.9997 on the pairs the spatial-block design can score, and a single unaided predictor, the count of cut edges on the baseline path, reaches 0.992. The unfitted rule built on that predictor needs no folds, applies to every pair, and catches every high-precision severance. Betweenness centrality is the weakest of six predictors, which is what should be expected on a network where cut edges are common rather than rare.

The identity of the critical links, however, does not survive the detector. The top eight links under the high-precision and bounding readings of the same imagery share no member. A structural severance model does transfer across inputs, at 0.715 to 1.000 when properly held out, but the targets are nested and the highest of those figures is close to mechanical, so we rest the recommendation on the unfitted screen rather than on the transfer. Structural fragility is a property of the network; which fragile pair a particular flood cut is a property of the sensor.

Two things follow that a reader should not take from the flood analysis alone. The first is that the same network holds three trunk links whose removal severs 812 settlement pairs each and splits the study area into two substantial pieces, and this flood, as this detector reads it, touches them on two acquisitions under one input out of four. A hazard-conditional ranking would not have found the largest structural exposure in the area. The second concerns the binary closure rule that produces the severance headline. Applying a published depth-disruption function removes severance entirely under the high-precision input at every depth swept, halves it under the bounding input at the depth that literature calls impassable, eliminates it at the depth a small car stalls, and raises the count of pairs merely delayed by up to a factor of 2.52. Where absorption fails is unchanged by the depth treatment, while whether the failure is a delay or a severance moves with the depth mapping. For operational use we recommend the structural screen, which needs no imagery, alongside any flood-specific ranking rather than in place of it.

References

Adelekan, Ibidun O. 2010. "Vulnerability of Poor Urban Coastal Communities to Flooding in Lagos, Nigeria." *Environment and Urbanization* 22 (2): 433–50. https://doi.org/10.1177/0956247810380141.

Aderamo, Adekunle J. 2012. "Urban Transportation Problems and Challenges in Nigeria: A Planner's View." *Prime Research on Education* 2 (3): 198–203.

Arrighi, C., M. Pregnolato, R. J. Dawson, and F. Castelli. 2019. "Preparedness Against Mobility Disruption by Floods." *Science of The Total Environment* 654: 1010–22. https://doi.org/10.1016/j.scitotenv.2018.11.191.

Bagloee, Saeed Asadi, Majid Sarvi, Brian Wolshon, and Vinayak Dixit. 2017. "Identifying Critical Disruption Scenarios and a Global Robustness Index Tailored to Real Life Road Networks." *Transportation Research Part E: Logistics and Transportation Review* 98: 60–81. https://doi.org/10.1016/j.tre.2016.12.003.

Balijepalli, Chandra, and Olivia Oppong. 2014. "Measuring Vulnerability of Road Network Considering the Extent of Serviceability of Critical Road Links in Urban Areas." *Journal of Transport Geography* 39: 145–55. https://doi.org/10.1016/j.jtrangeo.2014.06.025.

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

Berdica, Katja. 2002. "An Introduction to Road Vulnerability: What Has Been Done, Is Done and Should Be Done." *Transport Policy* 9 (2): 117–27. https://doi.org/10.1016/S0967-070X(02)00011-2.

Blanford, Justine I., Supriya Kumar, Wei Luo, and Alan M. MacEachren. 2012. "It's a Long, Long Walk: Accessibility to Hospitals, Maternity and Integrated Health Centers in Niger." *International Journal of Health Geographics* 11: 24. https://doi.org/10.1186/1476-072X-11-24.

Brandes, Ulrik. 2001. "A Faster Algorithm for Betweenness Centrality." *The Journal of Mathematical Sociology* 25 (2): 163–77. https://doi.org/10.1080/0022250X.2001.9990249.

Coles, Daniel, Dapeng Yu, Robert L. Wilby, Daniel Green, and Zara Herring. 2017. "Beyond 'Flood Hotspots': Modelling Emergency Service Accessibility During Flooding in York, UK." *Journal of Hydrology* 546: 419–36. https://doi.org/10.1016/j.jhydrol.2016.12.013.

Danczyk, Adam, Xuan Di, Henry X. Liu, and David M. Levinson. 2017. "Unexpected Versus Expected Network Disruption: Effects on Travel Behavior." *Transport Policy* 57: 68–78. https://doi.org/10.1016/j.tranpol.2017.02.002.

Duan, Yingying, and Feng Lu. 2014. "Robustness of City Road Networks at Different Granularities." *Physica A: Statistical Mechanics and Its Applications* 411: 21–34. https://doi.org/10.1016/j.physa.2014.05.073.

Echendu, Adaku Jane. 2020. "The Impact of Flooding on Nigeria's Sustainable Development Goals (SDGs)." *Ecosystem Health and Sustainability* 6 (1): 1791735. https://doi.org/10.1080/20964129.2020.1791735.

El-Rashidy, Rawia Ahmed, and Susan Grant-Muller. 2016. "The Evaluation of Redundancy for Road Traffic Networks." *Transport* 31 (4): 427–39. https://doi.org/10.3846/16484142.2016.1255913.

Freeman, Linton C. 1977. "A Set of Measures of Centrality Based on Betweenness." *Sociometry* 40 (1): 35–41. https://doi.org/10.2307/3033543.

Green, Daniel, Dapeng Yu, Ian Pattison, et al. 2017. "City-Scale Accessibility of Emergency Responders Operating During Flood Events." *Natural Hazards and Earth System Sciences* 17 (1): 1–16. https://doi.org/10.5194/nhess-17-1-2017.

Hopcroft, John, and Robert Tarjan. 1973. "Algorithm 447: Efficient Algorithms for Graph Manipulation." *Communications of the ACM* 16 (6): 372–78. https://doi.org/10.1145/362248.362272.

Jenelius, Erik. 2010a. "Redundancy Importance: Links as Rerouting Alternatives During Road Network Disruptions." *Procedia Engineering* 3: 129–37. https://doi.org/10.1016/j.proeng.2010.07.013.

Jenelius, Erik. 2010b. "User Inequity Implications of Road Network Vulnerability." *Journal of Transport and Land Use* 2 (3/4): 57–73. https://doi.org/10.5198/jtlu.v2i3.16.

Jenelius, Erik, and Lars-Göran Mattsson. 2012. "Road Network Vulnerability Analysis of Area-Covering Disruptions: A Grid-Based Approach with Case Study." *Transportation Research Part A: Policy and Practice* 46 (5): 746–60. https://doi.org/10.1016/j.tra.2012.02.003.

Jenelius, Erik, Tom Petersen, and Lars-Göran Mattsson. 2006. "Importance and Exposure in Road Network Vulnerability Analysis." *Transportation Research Part A: Policy and Practice* 40 (7): 537–60. https://doi.org/10.1016/j.tra.2005.11.003.

Kermanshah, Amirhassan, and Sybil Derrible. 2017. "Robustness of Road Systems to Extreme Flooding: Using Elements of GIS, Travel Demand, and Network Science." *Natural Hazards* 86 (1): 151–64. https://doi.org/10.1007/s11069-016-2678-1.

Kirkley, Alec, Hugo Barbosa, Marc Barthelemy, and Gourab Ghoshal. 2018. "From the Betweenness Centrality in Street Networks to Structural Invariants in Random Planar Graphs." *Nature Communications* 9 (1): 2501. https://doi.org/10.1038/s41467-018-04978-z.

Koks, E. E., J. Rozenberg, C. Zorn, et al. 2019. "A Global Multi-Hazard Risk Analysis of Road and Railway Infrastructure Assets." *Nature Communications* 10 (1): 2677. https://doi.org/10.1038/s41467-019-10442-3.

Latora, Vito, and Massimo Marchiori. 2001. "Efficient Behavior of Small-World Networks." *Physical Review Letters* 87 (19): 198701. https://doi.org/10.1103/PhysRevLett.87.198701.

Li, Daqing, Bowen Fu, Yunpeng Wang, et al. 2015. "Percolation Transition in Dynamical Traffic Network with Evolving Critical Bottlenecks." *Proceedings of the National Academy of Sciences* 112 (3): 669–72. https://doi.org/10.1073/pnas.1419185112.

Lint, J. W. C. van, Henk J. van Zuylen, and H. Tu. 2008. "Travel Time Unreliability on Freeways: Why Measures Based on Variance Tell Only Half the Story." *Transportation Research Part A: Policy and Practice* 42 (1): 258–77. https://doi.org/10.1016/j.tra.2007.08.008.

Makanga, Prestige Tatenda, Nadine Schuurman, Charfudin Sacoor, et al. 2017. "Seasonal Variation in Geographical Access to Maternal Health Services in Regions of Southern Mozambique." *International Journal of Health Geographics* 16: 1. https://doi.org/10.1186/s12942-016-0074-4.

Mattsson, Lars-Göran, and Erik Jenelius. 2015. "Vulnerability and Resilience of Transport Systems – a Discussion of Recent Research." *Transportation Research Part A: Policy and Practice* 81: 16–34. https://doi.org/10.1016/j.tra.2015.06.002.

Nkeki, Felix Ndidi, Philip John Henah, and Vincent Nduka Ojeh. 2013. "Geospatial Techniques for the Assessment and Analysis of Flood Risk Along the Niger-Benue Basin in Nigeria." *Journal of Geographic Information System* 5 (2): 123–35. https://doi.org/10.4236/jgis.2013.52013.

Nkwunonwo, U. C., M. Whitworth, and B. Baily. 2020. "A Review of the Current Status of Flood Modelling for Urban Flood Risk Management in the Developing Countries." *Scientific African* 7: e00269. https://doi.org/10.1016/j.sciaf.2020.e00269.

Ouma, Paul O., Joseph Maina, Pamela N. Thuranira, et al. 2018. "Access to Emergency Hospital Care Provided by the Public Sector in Sub-Saharan Africa in 2015: A Geocoded Inventory and Spatial Analysis." *The Lancet Global Health* 6 (3): e342–50. https://doi.org/10.1016/S2214-109X(17)30488-6.

Oyetubo, A. O., O. J. Afolabi, and O. A. Oluwaji. 2017. "Transportation Constraints to Rural Health Accessibility in Ogun Waterside Local Government Area of Ogun State Nigeria." *Journal of Research in National Development* 15 (1).

Papilloud, Tsolmongerel, and Margreth Keiler. 2021. "Vulnerability Patterns of Road Network to Extreme Floods Based on Accessibility Measures." *Transportation Research Part D: Transport and Environment* 100: 103045. https://doi.org/10.1016/j.trd.2021.103045.

Porta, Sergio, Paolo Crucitti, and Vito Latora. 2006. "The Network Analysis of Urban Streets: A Primal Approach." *Environment and Planning B: Planning and Design* 33 (5): 705–25. https://doi.org/10.1068/b32045.

Pregnolato, Maria, Alistair Ford, Vassilis Glenis, Sean Wilkinson, and Richard Dawson. 2017. "Impact of Climate Change on Disruption to Urban Transport Networks from Pluvial Flooding." *Journal of Infrastructure Systems* 23 (4): 04017015. https://doi.org/10.1061/(ASCE)IS.1943-555X.0000372.

Pregnolato, Maria, Alistair Ford, Sean M. Wilkinson, and Richard J. Dawson. 2017. "The Impact of Flooding on Road Transport: A Depth-Disruption Function." *Transportation Research Part D: Transport and Environment* 55: 67–81. https://doi.org/10.1016/j.trd.2017.06.020.

Scott, Darren M., David C. Novak, Lisa Aultman-Hall, and Feng Guo. 2006. "Network Robustness Index: A New Method for Identifying Critical Links and Evaluating the Performance of Transportation Networks." *Journal of Transport Geography* 14 (3): 215–27. https://doi.org/10.1016/j.jtrangeo.2005.10.003.

Sohn, Jungyul. 2006. "Evaluating the Significance of Highway Network Links Under the Flood Damage: An Accessibility Approach." *Transportation Research Part A: Policy and Practice* 40 (6): 491–506. https://doi.org/10.1016/j.tra.2005.08.006.

Starkey, Paul. 2016. "The Benefits and Challenges of Increasing Motorcycle Use for Rural Access." *Proceedings of the International Conference on Transport and Road Research (iTRARR)* (Mombasa, Kenya), March. https://research4cap.org/wp-content/uploads/ral/Starkey-ReCAPPMU_2016_BenefitsChallengesofIncreasingMotorcycleUseRuralAccess_iTRARR_160314.pdf.

Sullivan, J. L., D. C. Novak, L. Aultman-Hall, and D. M. Scott. 2010. "Identifying Critical Road Segments and Measuring System-Wide Robustness in Transportation Networks with Isolating Links: A Link-Based Capacity-Reduction Approach." *Transportation Research Part A: Policy and Practice* 44 (5): 323–36. https://doi.org/10.1016/j.tra.2010.02.003.

Suurballe, J. W. 1974. "Disjoint Paths in a Network." *Networks* 4 (2): 125–45. https://doi.org/10.1002/net.3230040204.

Tatem, Andrew J. 2017. "WorldPop, Open Data for Spatial Demography." *Scientific Data* 4:

170004. https://doi.org/10.1038/sdata.2017.4.

World Bank. 2019. *World: Measuring Rural Access, Update 2017/18*. No. ACS26526. World Bank. https://documents.worldbank.org/en/publication/documents-reports/documentdetail/543621569435525309/world-measuring-rural-access-update-2017-18.

Xu, Xiangdong, Anthony Chen, Sarawut Jansuwan, Chao Yang, and Seungkyu Ryu. 2018. "Transportation Network Redundancy: Complementary Measures and Computational Methods." *Transportation Research Part B: Methodological* 114: 68–85. https://doi.org/10.1016/j.trb.2018.05.014.

Xu, Xiangdong, Anthony Chen, Guangming Xu, Chao Yang, and William H. K. Lam. 2021. "Enhancing Network Resilience by Adding Redundancy to Road Networks." *Transportation Research Part E: Logistics and Transportation Review* 154: 102448. https://doi.org/10.1016/j.tre.2021.102448.

Yin, Jie, Dapeng Yu, Zhane Yin, Min Liu, and Qing He. 2016. "Evaluating the Impact and Risk of Pluvial Flash Flood on Intra-Urban Road Network: A Case Study in the City Center of Shanghai, China." *Journal of Hydrology* 537: 138–45. https://doi.org/10.1016/j.jhydrol.2016.03.037.

Zhang, Ning, and Alice Alipour. 2020. "Multi-Scale Robustness Model for Highway Networks Under Flood Events." *Transportation Research Part D: Transport and Environment* 83: 102281. https://doi.org/10.1016/j.trd.2020.102281.

Zhu, Shanjiang, and David Levinson. 2015. "Do People Use the Shortest Path? An Empirical Test of Wardrop's First Principle." *PLOS ONE* 10 (8): e0134322. https://doi.org/10.1371/journal.pone.0134322.

Zhu, Shanjiang, David Levinson, Henry X. Liu, and Kathleen Harder. 2010. "The Traffic and Behavioral Effects of the I-35W Mississippi River Bridge Collapse." *Transportation Research Part A: Policy and Practice* 44 (10): 771–84. https://doi.org/10.1016/j.tra.2010.07.001.